THE UNIVERSITY OF ADELAIDE

DEPARTMENT OF MECHANICAL ENGINEERING

The Automated Visual Inspection and Grading of Timber

submitted by

Peter J. M. Sobey, B.E.(hons)

for the Degree of Doctor of Philosophy

March 1989

THE AUTOMATED VISUAL INSPECTION AND GRADING OF TIMBER

PETER J. M. SOBEY

ERRATA

Spelling mistakes and word errors.

Note : B refers to the bottom of the page.

comment (page, line).

categories, categorise, (2.7, B-1), (6.1, 17), (6.1, B-1), (10.30, 15), (10.33, B-10), (10.35, 12), (10.35, B-8), (A3A.12, 6), (A3B.10, 6), (A4A.12, 6), (A4B.10, 6).

in millimetres, (A3A.12, 9), (A3A.10, 9), (A4A.12, 9), (A4B.10, 9).

exception, (2.2, 9), (9.33, 12).

are then explored, (3.1, 8).

Paul (1988), (3.11, 8).

minima > minimum, (3.12, 2), (3.12, 4).

minimas > minima, (3.12, 4).

"flatness", (3.20, B-3).

to be, (4.2, 1).

computational, (4.2, B).

position, (5.6, 2).

digitised, (5.7, 9).

principles, (5.11, 4).

separate, separated, (5.13, B-6), (5.14, B-10), (10.29, B-2).

computation, (6.1, 10).

to further examine the issue of, (6.6, 9).

Starting, (6.8, B-5).

dropping, (6.11, 1).

criteria, (7.4, 5).

In equation 7.2 the two terms on the right hand side should be added, not subtracted, (7.7).

In equation 7.5 the term in brackets should be squared, (7.8).

In equation 7.6 the denominator should be $\bar{\sigma}_k$, (7.8).

labelled, (7.10, B-6).

means that a deterministic method can be used., (7.10, B-5).

It is preferable to have a cost function with only a single minimum to avoid getting stuck at a local minimum., (7.11, B-4).

conventional notation, (Ho and Agrawala (1968))., (7.12, 16).

square-law function, (7.14, 3), (8.2, 20), (8.2, 21).

position, (7.15, B-1).

figure 7.9, (7.19, 3).
normalisation, (8.5, B-4).
mapped onto the 16 grey levels, (8.7, 13).
homogeneity, (8.8, 5), (8.10, B-7), (8.10, B-1), (8.20, B-1), (9.38, 12), (9.42, 16), (9.47, 6), (9.47, B-2), (9.53, B-7).
describes, (8.9, B-7).
than, (8.16, B-3).
is established, (9.1, B-4).
STABILITY, (9.5, B-6).
determines the amount, (9.18, 2).
boundary, (9.18, 3).
artefacts, (9.33, 20).
determining, (10.27, B-4).
is divided, (10.34, 1).
are presented, (10.34, B-3).
excludes, (11.4, B-2).
several, (11.7, B-2).
In table 12.1 (12.5) and table 12.2 (12.6) the terms 190*19 and 42*19 refer to the size of the timber section in millimetres.
authors', (12.10, 1).
are investigated, (13.2, 20).

Points of Clarification.

Page 6.4 (last paragraph) over to page 6.5 explains the measure kurtosis of the grey level histogram. "Flatness" refers to the appearance of the image. A uniform image will appear "flat" to an observer and will have a histogram with values clustered about the mean. Not many pixels wil have grey levels that differ markedly from the mean. An "unflat" image will not appear of uniform intensity to an observer and the histogram will be characterised by pixels far from the mean.

Page 6.4, line B-9, seeks to visualise the measure of skewness. This sentence, however, is incomplete in its explanation and should refer to the fact that if the histogram is symmetric about the mean the skewness will be zero.

Page 6.10. The first paragraph can be continued in the following manner:

... No meaningful threshold can be deduced from the skewness data alone due to the nature of the distribution. Although clear areas consistently have a value of skewness close to zero, areas that contain features have skewness values that can vary from large negative values to large positive values. A skewness threshold only extracts a small number of features at any one value and so is not useful in this capacity.

Page 7.5, second paragraph. The coefficient referred to is that which corresponds to the measure in question in the vector that defines the decision boundary. Obviously, if the coefficient is very small compared to the other coefficients of the decision boundary vector then the value of that measure in a sample will have very little influence on the position of the sample with respect to the decision boundary.

Page 7.12, third paragraph. A brief reference is made to augmented and unaugmented pattern space. This is explained briefly earlier (equation 7.7) and a full description can be found in Meisel (1972), page 63.

The use of a linear decision boundary to divide pattern space into two classes, as described in chapter 7, is a specific case of the more general perceptron training method. In particular it is a single layer perceptron with a single threshold element. This was felt to be sufficient for the task at hand.

Figures 7.4 to 7.12 are illustrations of all the samples of the training set and show the relation of samples in pattern space during an iteration (figures 7.4 to 7.9).

In section 12.2 (12.4) the fact that the camera pixels are not square is not of major concern. The resolution of features across the board is required to a greater degree of precision than that along the board. This can be adjusted with a linear array camera by varying the exposure time for each line grabbed. The resolution across the board is defined by the maximum board width and the number of pixels in the CCD array.

SUMMARY

A method has been developed to visually detect and discriminate the features on the surface of sawn planks of radiata pine for the purpose of grading the timber. The features include small and large knots, cone holes, pith, bark, wane, resin and holes. The grade of the plank is determined by the number, distribution and type of features present. The automation of the visual inspection of timber for grading to a national standard has not previously been attempted which makes the present study unique.

It is postulated that the grey level histogram of the image of the wood is sufficient to discriminate the features from the background of prominent growth rings. The image of the wood is digitised from a charge coupled device (CCD) camera and segmented into smaller sub-images, called local areas. Statistics of the distribution of pixel intensities within each local area place it in a feature-measure space. A decision boundary is derived from a labelled set of samples that divides the feature-measure space into two half-spaces: one containing local areas that enclose a part of a feature, and the other containing local areas that contain only growth rings. This method allows over 95% of the local areas to be correctly classified.

Local areas that contain parts of the same feature are identified and merged together to form feature areas. The grey level histogram of the feature area is postulated as being sufficient to separate and size the features within a feature area. It is further postulated that an identification of the type of feature can be made with shape measures. This is supported by the analysis of a series of images containing a range of features which reveals the strengths and weaknesses of the methods used to discriminate them.

A comprehensive study of the use of texture measures in the detection and discrimination of features concludes that these measures are required to improve the performance of the algorithms based on histogram methods to a level that will allow commercialisation of an automated inspection system. The design of a production system is described and found to be a realisable project.

TABLE OF CONTENTS

STATEMENT OF ORIGINALITY

The material contained in this thesis is original and has not been submitted or accepted for the award of a degree or diploma at any other university. To the best of the author's knowledge and belief, no material which has previously been published or written by another person is included, except when due reference is made in the text.

Peter J. M. Sobey
March 1989

PERMISSION TO COPY

The author consents to the thesis being made available for loan and photocopying provided that it is accepted for the award of the degree.

Peter J. M. Sobey
March 1989

ACKNOWLEDGEMENTS

I would like to express my gratitude to the many people who were a part of this project. Particular thanks are due to my supervisor Euan Semple for letting me follow my interests and the guidance he provided toward a practical application. Without the opportunities provided with his help and his timely comments I would not have come so far.

The work in this thesis was carried out in the Department of Mechanical Engineering of the University of Adelaide under the aegis of Professor R. E. Luxton. I would like to thank Professor Luxton for his interest in this work as a new area of study in the Department and for the opportunity to conduct this research.

To the Departmental staff I would like to thank Silvio De Ieso and Alan Mittler in particular for their assistance with electronics projects. Thanks also to the other workshop and office staff for their friendly assistance whenever it was requested.

My fellow post-graduate students deserve thanks for their companionship during my years of study, especially Andrea Bunting for her comradeship.

Video equipment used in this project was borrowed from the Woods and Forests Department of South Australia. Special thaks are due to the Director, Roger White, who saw merit in the project and made the equipment available. The wood used in the study was video-taped at the Mount Gambier Mill of the Woods and Forests Department. This was made possible with the friendly and interested co-operation of the staff at Mount Gambier, in particular Chris Jones, Graham Bayly and Rick Garret of the Quality Control team. Many thanks also to Les Cunningham and Bob Clintburg for their interest and co-operation.

Thanks are due to my parents for their support and encouragement throughout my university days, to Lotus Cavagnino for her patience and help in the preparation of material, and Ishtar Cavagnino for giving me reason to persevere.

Finally, I acknowledge the financial support in the form of a Post-graduate Research Scholarship from the University of Adelaide.

NOMENCLATURE

AUSTRALIAN STANDARD SPECIFICATION

The grading rules as specified by the Standards Association of Australia.

BACKGROUND

The wood is considered to consist of the features, knots, etc, and background, which is clear wood containing only growth rings.

BACKGROUND AREA

That area within an image or sub-image that is not part of a feature.

BLOB

Feature pixels within a thresholded area that are all connected by the four neighbour rule.

BLOB MEASURE

A measurement of a blob used to describe a particular aspect of the blob.

CLASSIFICATION ALGORITHM

An algorithm that is used to classify a local area into the two classes of feature and clear.

CO-OCCURRENCE MATRIX

Spatial grey level dependence matrix, SGLDM. Used in the analysis of texture.

DECISION BOUNDARY

A hyperplane in pattern space that defines the two areas of feature and clear within the space.

EQUAL PROBABILITY QUANTISING

A means of reducing the number of grey levels in an image while retaining texture information.

FEATURE

Any aspect of the wood surface that is distinct from the pattern of growth rings. Includes knots, pith, resin, bark, wane, holes and needle trace. A description of each of these features is found in section 2.3.

FEATURE AREA

A rectangular region of an image that is identified as containing one or more features. Feature discrimination processing operates on the feature area level.

FEATURE DETECTION

That part of the inspection process concerned with the detection of the presence of features within local areas.

FEATURE DISCRIMINATION

That part of the inspection process concerned with the determination of the size, shape and type of feature(s) within a feature area.

FEATURE EXTENT

The location and size of the smallest rectangle that encloses a feature blob.

FEATURE MEASURE

A measure used to identify a feature area apart from a clear local area. Tonal and texture measures are investigated in this study.

FEATURE LOCAL AREA

A local area that contains a portion of a feature.

FEATURE SIZE

The size of the feature as specified by the Australian Standard Specifications and by experienced timber graders.

GRADING

The application of the grading rules of the Australian Standard Specification to radiata pine. The planks are placed into one of a number of grades depending on the size, type and distribution of features on the surface of the wood.

GRADING RULES

The rules that define the type, size and distribution of features that determine a certain grade of timber. Specified by

the Standards Association of Australia in a number of standards for Radiata Pine. Those for Sawn Boards are found in Appendix 5.

GREY LEVEL HISTOGRAM

See histogram.

HISTOGRAM

An array that lists the number of pixels of each grey level in the (sub-)image. Because there are 256 grey levels the histogram has 256 elements. Normalisation is achieved by dividing each histogram value by the total number of pixels in the image. The sum of the elements of the normalised histogram array is 1.

IMAGE

The digital representation of a portion of plank surface stored in computer memory.

LINEAR HYPERPLANE

A linear function that divides a pattern space into two parts. A straight line is a 1-d function that divides a 2-d plane into two parts; a 2-d flat surface is a linear function that divdes 3-d space into two parts; a (n-1)-d linear surface divides a n-d space into two parts. Synonymous with a linear decision boundary.

LOCAL AREA

A sub-image of fixed size. A size of 64 * 64 pixels is used in this study unless stated otherwise. Local areas do not overlap.

PATTERN CLASSIFICATION

The classification of a sample by the position of the sample in pattern space. The position of the decision boundary determines the class of a sample.

PATTERN SPACE

A multi-dimensional space with each dimension representing a measure of the local area. Calculation of the feature measures of a local area defines a vector, $\mathbf{x}$, that represents the local area in this pattern space.

PIXEL

A picture element. A digital image is an array of pixels, each pixel representing the intensity of light at that position in the image as a number from 0 (black) to 255 (white)

RADIATA PINE

A species of pine that originated from California. Used as a plantation timber in southern Australia because of its rapid growth.

RPAA

Radiata Pine Association of Australia. An industry association.

SGLDM

Spatial grey level dependence matrix. Used in the analysis of texture.

TEXTURE MEASURES

Measures of a digital image that quantify aspects of the spatial and grey level relationships between pixels in an area. A description of all the texture measures used in this study is found in section 8.1.2.

THRESHOLD

A value that determines the transition between one state and another. A threshold grey level is used to create a blob image of only 2 grey levels. All pixels in the area with a grey level above the threshold value are classed as background and given the value 255 (white). All those equal to or below the threshold value are classed as feature and given the value 0 (black). The result is a black feature blob on a white background.

TONAL MEASURES

Measures of a digital image that quantify the distribution of grey levels. Based upon the grey level histogram.

WEIGHT VECTOR

The vector, w, that describes the position of the decision boundary in pattern space.

LIST OF SYMBOLS

ASM	angular second moment texture measure
CON	contrast texture measure
COR	correlation texture measure
CSH	cluster shade texture measure
CPR	cluster prominence texture measure
$C(\mathbf{x})$	decision function
dz	dead zone
ENT	entropy texture measure
$L[C(\mathbf{x}),i]$	loss function
lfn	loss function parameter
LHOM	local homogeniety texture measure
lossrat	loss ratio
M	number of samples in the set
n	number of feature measures
N	number of classes
N_g	number of grey levels
$P(\theta,d)$	summation Spatial Grey Level Dependence Matrix
$\tilde{R}(\mathbf{w})$	cost function
$\tilde{R}'(\mathbf{w})$	approximation of cost function
S	set of all samples, $S_0 + S_1$
S_0	set of clear samples
S_1	set of feature samples
$s_\theta(i,j\|d)$	(i,j)th element of $S_\theta(d)$
$S_\theta(d)$	normalised Spatial Grey Level Dependence Matrix
tdv	texture displacement vector, alternatively referred to as the inter-sample spacing, d
v	gradient function of the cost function
w	weight vector
x	sample vector in pattern space

$\bar{x}_k$	mean of sample measure, k
ϵ	gradient increment
$\Phi_{(d,\theta)}$	joint conditional probability density function
λ	step
$\bar{\sigma}^2_k$	variance of sample measure, k
ω	$(w_1, w_2, \ldots, w_n)$

1. INTRODUCTION

Australia has large areas of pine forest to satisfy the local market demand for softwood (figure 1.1). These forests are a monoculture of the Californian species Pinus Radiata (once known as Pinus insignis and later as Remarkable Pine on account of its remarkable rate of growth). Most of this timber is mechanically stress graded for structural purposes but a significant proportion is being visually graded where the

Figure 1.1 Kuitpo forest south of Adelaide.

application is more decorative than structural, e.g. the panelling of walls. Pine boards that are to be graded by appearance are graded to rules established by the Radiata Pine Association of Australia and the Standards Association of

Australia. The grade is decided by the frequency, distribution, size, and type of the visual features on the surface.

1.1 THE NEED FOR AUTOMATED VISUAL INSPECTION

The grading is done at present by relying on the skill and alertness of human inspectors and this is the source of a quality control problem. Human judgement is difficult to sustain, particularly at speed and with long boards, and inspectors suffer from fatigue and boredom even though the job is rotated. These factors combine to give a variable result. When presented with a board that falls between two grades, the inspector tends to downgrade it, and this leads to an undervalued product. Usually about 6% of the wood is out of grade and on occasions this can rise to over 24%. Mills presently in operation produce a 6 metre board every 3 seconds (2 m/s). Those that are proposed for the 1990's have speeds in excess of 3 m/s, and a limitation to the introduction of these faster machines is the visual inspection process. The technology to automate visual inspection has only become commercially available over the last few years and no automated system for wood grading has yet been implemented successfully.

Automating this inspection process has received motivation from a number of areas. The first of these is the degree of variability inherent in the current inspection process due to fatigue, boredom, and differences in operator experience. There is a need to improve the quality of the grading process to maximise return on the product. The second is the increase in the mill speed: operators are finding it increasingly difficult to grade the wood effectively at the faster rates. The third is the need to maximise productivity with the minimum of capital expenditure. Is an automated inspection process more cost effective than a duplication of the grading lines?

The variability of the present system of grading is revealed by examining the results of quality control inspections of graded packs (this is presented in more detail in the following chapter, tables 2.4, 2.5, and 2.6). The average number

of pieces from a pack of standard grade that are below grade is 5.9%. This is mainly due to excessive cone holes. A higher standard of quality is demanded for furniture grade timber and this showed an average of 17.9% below grade. The higher grade has a greater proportion of low grade timber within it. The lowest grade of timber is called merchantable grade and this presented an average of 24.0% above grade. A majority of this wood was either standard grade or could be docked to give over half the length as standard grade. These figures were obtained by examining the quality control officer's reports for 25 pack checks. In this procedure a pack of timber that is packaged and ready for dispatch to market is opened and regraded by the quality control officer and the results recorded indicating the number of pieces of timber wrongly graded. The sample totalled 3255 pieces of timber. Recent efforts have been made to improve the level of quality control by a more rigorous training of the graders but the variability factor is well illustrated.

	STANDARD	MERCHANTABLE	TOTAL
VOLUME (m^3)	42,750	14,250	57,000
PERCENTAGE (%)	75.0	25.0	100.0
VALUE ($M)	19.238	3.562	22.800

Table 1.1 Typical breakdown for a mill with an input of 150,000 cubic metres of logs per year. Value of sawn timber is $450 and $250 per cubic metre for standard and merchantable grades respectively.

The costs of this level of grading error can be estimated as in table 1.1. Consider a mill with an annual log intake of 150,000 cubic metres (the annual sawlog usage of the Green Triangle region of South Australia and Victoria is 700,000 m^3 (Roughana (1988))). A typical value for dry, dressed recovery is 38%, the remaining 62% being comprised of sawdust, shavings and water. Of this 57,000 cubic metres of saleable timber a

typical breakdown is 75% standard grade and 25% merchantable grade. If the price for standard grade is taken to be $450 per cubic metre and that for merchantable grade to be $250 per cubic metre, the output has a sale value of $22.8 million.

It is known that 24% of the merchantable grade timber can be placed in standard grade if the inspection is performed precisely. If an improved grading system reduces this to 5% out-of-grade, that is 19% of the merchantable grade timber is regraded as standard grade then the output will have a saleable value of $23.341 million, an increase in return of $541,000 (table 1.2).

	STANDARD	MERCHANTABLE	TOTAL
VOLUME (m^3)	42,750	14,250	57,000
CHANGE (m^3)	+2,707	-2,707	
VOLUME (m^3)	45,457	11,543	57,000
PERCENTAGE (%)	79.75	20.25	100.0
VALUE ($M)	20.456	2.886	23.341
VALUE CHANGE ($)			+$541,000

Table 1.2 The increase in value when 19% of the merchantable grade timber, that is currently wrongly classed, is upgraded to standard grade.

However, of this up-graded timber only one third can be put straight into standard grade. The other two thirds of the timber can only be put into this grade if it is docked to a half length or greater. If it is assumed that only one half of this two thirds is placed in standard grade then two thirds of the timber removed from merchantable grade is placed into standard grade. A smaller amount of material is moved into standard grade than in the previous situation. Table 1.3 indicates that this reduces the value of the gain considerably.

	STANDARD	MERCHANTABLE	TOTAL
VOLUME (m^3)	42,750	14,250	57,000
CHANGE (m^3)	+1,805	-2,707	
VOLUME (m^3)	44,555	11,543	56,098
PERCENTAGE (%)	79.42	20.58	100.0
VALUE ($M)	20.050	2.886	22.935
VALUE CHANGE ($)			+$135,000

Table 1.3 19% of merchantable grade material, wrongly graded, is removed but only two thirds of this is placed into standard grade.

If a perfect grading system graded the timber 100% true to grade and again only two thirds of the merchantable grade timber is upgraded to standard grade then the results show a loss. The amount of wrongly graded material removed from standard grade reduces the final value more than the material going from merchantable grade to standard increases it. Table 1.4 gives the figures applied to the original example with 5.9% out of grade in standard grade and 24% out of grade in merchantable grade of which two thirds finds its way into standard grade.

Clearly, perfect grading is not in the economic interest of the timber processor. Ideally, the timber processor should realise the maximum return while conforming to the Australian Standards that allow up to 5% out of grade material. This is the situation where the standard grade contains 5% of material from the lower grade and the merchantable grade contains no material from the higher grade.

	STANDARD	MERCHANTABLE	TOTAL
VOLUME (m^3)	42,750	14,250	57,000
CHANGE (m^3)	-2,522	+2,522	
CHANGE (m^3)	+2,280	-3,420	
VOLUME (m^3)	42,508	13,352	55,860
PERCENTAGE (%)	76.10	23.90	100.0
VALUE ($M)	19.129	3.338	22.467
VALUE CHANGE ($)			-$333,000

Table 1.4 Graded 100% true to grade. 5.9% is removed from standard grade and placed in merchantable grade. 24% is removed from merchantable grade and two thirds of this is placed in standard grade.

If the standard grade did not contain any merchantable grade material then a 5% price increase would increase the value of the standard grade timber by $956,000 and turn the loss into a gain of $623,000. This extra value would not be realisable in the current market situation, but this is not the whole issue.

Litigation is a growth industry in the United States of America and, although it is of a lower level of incidence in Australia, it can have devastating consequences on the financial position of the defendant of a test case. If a litigation case proceeded against a timber processor on the grounds that the claimant suffered loss as a result of timber being out of grade the defendant could not rely on the Standard specifications as a defence. The defendant would have to prove that every effort had been made to conform to the Standards and would deflect the prosecution by pointing out that there is a requirement for the buyer to exercise responsibility in the application of a product with a natural variability such as wood. Whatever the outcome of the case, timber processors would have to ensure that they employed the best techniques to conform to the Standards that

were available. The processor with the ability to provide the closest adherence to the Standards would then be in a position to exploit the market situation to the best advantage.

While the preceeding calculations highlight the costs involved in wrongly grading timber, the point must be made that it is playing with hypothetical figures. The real benefit of an automated inspection system is the increase in control that the operator gains over the production process. This control will enable a more consistently graded product to be produced with a reduced labour requirement. By reducing the number of graders and dockers on an inspection line by 8 the savings are of the order of $280,000 per shift per year. A more consistently graded product enables end-users, such as furniture manufacturers, to operate within tighter economic constaints secure in the knowledge that they will not receive a poor quality batch of wood that can destroy their profit margin.

A common problem with the introduction of new technology is that it must fit into an already existing production environment. The mills at Mount Gambier are a case in point. New milling equipment to replace the present machines in the 1990s promises an increased throughput of 50%. This places an increased burden upon the teams of graders that must keep pace with the output. A simple solution would be to duplicate the grading facilities and staff to cope with planned and future throughput but the fact is that there is a limited amount of floor space for this to be implemented. The capital cost involved in such a solution has prompted the consideration of other methods, namely automated visual inspection, that also hold the promise of increasing the efficiency of the grading process.

The cost factor, as well as the space factor, already inhibit measures that would lead to an increase in productivity. The quantity of merchantable grade timber that can be up-graded by docking is significant and this can be accomplished using existing technology. The cost, however, is prohibitive at present and there is insufficient floor space to install the docking machines. It is envisaged that an automated grading system would

incorporate an automated docking facility and so eliminate the pressure on floor space.

From the levels of timber that are out of grade it can be seen that there is a great deal of room for improvement. Standard grade timber is being undervalued as merchantable grade and furniture grade timber contains an unacceptable level of lower grade wood. The former contributes directly to a loss of return and the latter erodes consumer confidence in the value of the timber as a whole.

Automation of the visual inspection process is necessary to improve the quality control of the product by eliminating the variability of the human inspection. This will increase productivity and remove the need for more floor space with the introduction of faster planing machines. Improved quality control of the grading will maximise returns and increase consumer confidence.

1.2 PREVIOUS WORK

Previous studies into the automated inspection of timber have concentrated upon the determination of cutting strategies. A round log is cut into rectangular section planks. The planks that extend to the edge of the log do not have a rectangular section. The absence of a square edge due to the outside edge of the log is called wane and sawing strategies are methods that remove the wane and leave the maximum section of plank. Laser scanning systems detect the wane and position the saws to optimise the return. More sophisticated systems also incorporate a degree of knot detection to direct docking cuts. These types of system are implemented in European timber mills where the output is mainly undressed boards that are classed into only one, or at most two, grades. This timber is bought by manufacturers who dress the timber for their own purposes.

Australia appears to be unique in that the timber mills plane nearly all their production for sale direct to retailers and manufacturers. This finished timber is also graded by the producer in the finishing mill to a number of Australian

standards which is also unique. Structural timber is generally mechanically stress graded and appearance grade timber is visually graded. No study has previously been conducted into the feasibility of automating the visual grading of appearance grade timber.

1.3 OBJECTIVES

The objectives of this study are as follows:

1 - Investigate the aspects of the visual appearance of timber that distinguish features as distinct from clear areas of board.

2 - Quantify the ability of combinations of tonal and textural measures to detect features. Show that these measures distinguish particular types of areas.

3 - Consider methods for determining the size of features given that the edge of a feature is not always clearly defined. A labelled set of sample images is used to verify the results of the automated process.

4 - Estimate the feasibility of constructing a system that can perform the inspection task at the speed of timber production.

1.4 LAYOUT OF THIS THESIS

Chapter 2 begins with a description of the production of timber. The origin of the various types of features present in the final product is discussed, and the problem of visual grading is examined.

Chapter 3 is a literature survey that explores the various approaches to automated visual inspection. This reveals that the problem of grading timber to a national standard is one which has not previously been addressed.

The approach to the inspection process is outlined in chapter 4 and a description of the equipment used in the study is presented in chapter 5.

Chapter 6 presents an initial investigation into the applicability of tonal measures to discriminate feature areas and the optimum size of segmentation of the image.

The theory of pattern classification is developed in chapter 7 and a visual representation is presented. This method of combining feature measures is applied with the introduction of texture measures in chapter 8.

A more comprehensive examination of the classification algorithm is applied to a set of images in chapter 9 and the influence of the tonal and texture measures investigated.

Two methods of determining the size of features are developed in chapter 10 and the results presented in the appendices.

Chapter 11 demonstrates the ease of codifying the rules for grading timber and chapter 12 presents proposals for future work toward the design and construction of a production system that automate the grading of radiata pine.

Conclusions are presented at the end of each chapter where appropriate and summarised in chapter 13.

2. THE TIMBER INDUSTRY

2.0 INTRODUCTION

Radiata pine is a species native to California and the first experimental plantings in South Australia commenced in 1876. It showed exceptional growth compared to other species tested which included native eucalypts. Regular plantings

COUNTRY	AREA ('000 ha)	PROPORTION OF GLOBAL ESTATE
NATIVE STANDS		
USA (Californian mainland)	7	<0.25%
MEXICO (islands)	0.15	
EXOTIC PLANTATIONS		
NEW ZEALAND	1,010	33.1%
AUSTRALIA	600	19.7%
CHILE	1,090	35.7%
SPAIN	270	8.9%
OTHER (mainly SOUTH AFRICA)	75	2.5%
WORLD TOTAL	3,050	

Table 2.1 Estimated total global estate of Pinus Radiata as at April 1986 (from Lavery (1986)).

commenced in 1907 and by 1910 over a million trees were growing on 800 hectares. Commercial plantings began in Victoria in 1917, New South Wales and Tasmania in 1921, Western Australia in 1922, Australian Capital Territory in 1925, and Queensland in 1927. The total area planted to radiata pine in Australia in March 1985 was

574,000 hectares (Lavery (1986)) and new plantings cover about 28,000 hectares per annum. Table 2.1 shows that Australia grows about 20% of the worlds radiata pine and also shows the relative contributions of the other major growers which are Chile, New Zealand, South Africa, Spain, Italy, India, and Fiji.

Radiata pine contributes most of the softwood and nearly 25% of all timber produced in Australia. The area planted in each state is shown in table 2.2. This provides nearly all the local demand for softwood with the exeption of Queensland. The reason for this is that the Queensland climate is unsuitable and they also have access to native species (Hoop Pine).

STATE	AREA (ha)	PROPORTION (%)
NSW	194,232	33.8
Vic.	179,479	31.2
Qld.	3,185	0.6
WA	44,127	7.7
SA	78,342	13.6
Tas.	61,362	10.7
ACT	13,688	2.4
TOTAL	574,415	

Table 2.2 Confirmed plantation area statement of radiata pine in Australia as at March 1985 (Lavery (1986)).

2.1 THE FOREST

The plantations of radiata pine in South Australia are grown on former farmland as this is usually well watered land in this the driest state. Because the timber industry must compete with farmers for land it is not expected that the area under plantation in SA will increase significantly in the future. The industry in this state has therefore reached a situation where it

is in a steady cycle of harvesting and planting at a continuous sustainable level (although events such as the Ash Wednesday bushfires of 1983 have put hiccups into the system).

Rainfall requirements for radiata pine are in the range of 660mm to 890mm per annum in areas with wet winters and dry summers and the trees will grow from sea level to an altitude of 915 metres. The growth rate of the timber is truly remarkable. In its native California yields are as high as 5.4 cubic metres per hectare per year, based on a 25 year rotation (Scott (1960)). Spain has been able to improve this up to 7.5 cubic metres per hectare but in Australia yields are as high as 24.5 cubic metres per hectare. Growth rates of one sixth this value are considered good among softwood plantations in Europe.

Figure 2.1 A plantation of radiata pine showing the regular spacing and uniform height.

The trees are planted in compartments of 15 to 20 hectares being uniform in age and spacing (figure 2.1). This induces uniform upward growth of the limbs since no tree receives

light from the side. Exceptions to this are the trees that form the border of the stand along firebreaks and access trails. These break trees receive a proportion of their light from the side and develop significant side branches along their entire length (figure 2.2). A 20 hectare square compartment of middle aged trees will contain about 6,400 trees at a spacing of 5.6 metres. Of these 320 will be break trees which constitute 5% of the compartment. The pattern of growth of the branches has important implications for the quality of timber produced. Large side branches will result in unacceptably large knots in the sawn boards which have to be docked to remove them. This results in shorter usable lengths from these trees.

Figure 2.2 A tree growing at the edge of a stand. Note the significant side branches and the regular spacing of the limbs.

As the forest grows it is thinned to optimise the spacing of the trees. The trees in a compartment are initially planted at a spacing of 2.44 metres to give a tree density of 1680 per hectare. Wider spacing (3.5 metres) leads to an increase in the incidence of forking of the leader stem and competition by lateral branches. These malformed stems are more easily removed out by thinning a closer spaced compartment. A first thinning is performed at 8-12 years and is usually left on the ground.

Commercial thinnings are conducted at 15-20 years and every 5 years from then on to cull the poorer trees and to reduce competition for resources for the remainder. These thinnings are used for poles or chips or are pulped for paper products. The final forest grows to a height of around 40 metres at a density of 160 trees per hectare (7.9 metre spacing). Figure 2.3 shows the development of a managed plantation forest of radiata pine.

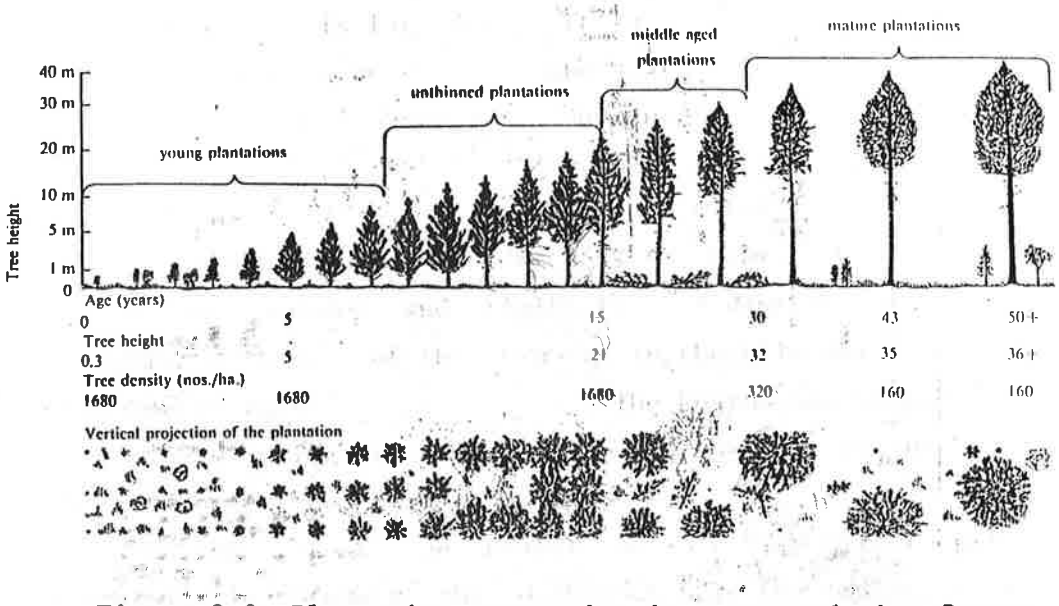

Figure 2.3 Plantation succession in a managed pine forest. (From B. C. Gepp (1986))

Radiata pine is considered to be a self tending forest in that the lower branches fall off as they die. Branches that do not fall off are encased by the growing tree and form unsound knots in sawn timber. It is not considered economic by the SA Woods and Forests Department to prune these lower branches except for a small quantity of surface peeler logs (these are logs that are "peeled" to form the surface sheets for plywood and so must be of a clear quality, not a requirement for the other layers of the ply). Queensland, WA, and New Zealand are known to prune the lower branches to increase the amount of clear timber that can be extracted. By low and high pruning, the knots can be confined to a central core 100mm in diameter. Therefore, a 300mm diameter log with a 100mm diameter knotty core would have 89% of the total

volume knot free. This practice leads to a greater proportion of higher value, clear timber but must be balanced by the fact that for appearance grade the knots are not necessarily defects to be eliminated but are features that enhance the appearance. Furthermore, the cost of pruning is significant and is committed relatively early in the life of the plantation. Consequently its cost including interest at harvest must be offset by a certain higher value per m^3.

2.2 THE TIMBER MILL

The forest is logged when it is about 45 - 50 years old. The logs are debarked and sawn into boards called green timber. The cutting pattern of the log is designed to create boards that utilise the anisotropic nature of the wood to best advantage. The green boards are stacked and placed into an kiln with controlled temperature and humidity to be dried. This stabilises the wood to a 15% moisture content and limits further dimensional changes. The stacking technique of the timber in the kiln is designed to put the boards into a straight set. The boards are unloaded for entry to the dressing mill that planes the surfaces of the boards to the finish size and surface quality.

At the exit of the planer an inspector directs planks with oversize features to the dockers where the feature is cut out and the two shorter pieces are then redirected for grading if of sufficient length. Those boards that do not require docking pass along a conveyor to four pairs of inspectors, one at either end of the six metre planks. At a maximum range of three metres the inspector must detect and size features to within 2 millimetres and apply the RPAA grading rules that restrict the proximity and number of features within each plank. The grader places the plank into one of the catagories of clear, joinery, select, standard and merchantable grades. Sometimes only two or three grades are used depending on the the type of wood being graded. For instance, flooring is graded into the catagories of standard grade, for floors that require a good appearance, and

merchantable grade, for floors that are intended to be covered by carpet.

The areas of difficulty in human visual inspection are the rate of inspection, the resolution required, and the repetitive nature of the task. The human inspectors are at present working at a maximum rate beyond which quality would suffer greatly. The only way to keep up with a faster throughput of material is to increase the number of graders at the output of each mill. This is not a desirable option as there will be difficulties in material handling with production then needed to be divided into separate grading lines each running within the capacity of a manual grader. The RPAA grading rules require the size of features to be determined to +/- 1mm. The ability of a person with good eyesight to be able to judge whether a bark inclusion is 5mm wide and in grade or 7mm wide and out of grade from three metres within a couple of seconds stretches credibility. The ability to be able to make these decisions reliably and constantly for hours on end taxes human abilities to the extreme.

2.3 THE FEATURES OF RADIATA PINE

Radiata pine is a knotty timber and it is the presence of these knots that make it visually appealing. For this reason the knots are referred to as features, not defects, and the quality, size and distribution of these and other features is used to determine the grade of the timber. These features can be listed as follows:

> Knot - caused by a branch growing from the side of the tree. Knots vary in size from less than one millimetre in diameter to over 35mm. They appear as a dark oval (the exact size depending on the angle at which it is cut) with the colour being darker as the branch gets older and dies due to an increasing resin content. When the knot first starts to grow it is very faint, appearing as little more than a

distortion of the growth rings. Knots can be sound if they are still intergrown with the bulk of the tree. When the branch dies it no longer grows with the tree and becomes encased. An encased knot has no structural integrity with the tree. Between these two extremes are partially encased knots.

Spike knot - A branch that is close to the centre of the log will tend to be cut longitudinally rather than laterally. This gives a long feature that is often seen to radiate from pith. Timber cut from all but the centre of the log is cut in such a way that any branch is cut laterally. Spike knots are only associated with cuts through the centre of the log which are usually the wider section planks.

Hole - A hole is the result of an encased knot that has fallen out. This is a serious defect in an appearance grade timber.

Pith - The soft central core of the log. In sawn timber it is recognised as a dark band on the surface that can occasionally be more than 10mm wide.

Resin and Bark - Bark only appears in the interior of a log if it becomes entrapped. This usually occurs at a branch, either as encasement of a dead limb, or in the angle of a steep branch. In these cases it is associated with a knot. Occasionally it can become encased in a longitudinal band between growth rings. In either case it represents a structural discontinuity and is limited by its width. Resin appears between growth rings as a dark band. Is not as dark as bark and does not always create a loss of integrity.

Blue Stain - A discolouration of the timber due to fungi.

Chipping - Cracks and voids in the centre of a knot caused because the knot is cut across the end of its grain while the rest of the plank is cut along the grain.

Cracks and Splits - Separation along the grain due to differential shrinkage during drying.

Wane - Bark (or lack of wood from any cause) on the surface of the plank. Found in boards cut from near the edge of the log.

Hit and Miss - Areas of the plank that have failed to receive a smooth finish from the milling cutter.

Bow, twist, spring and cupping - These describe gross deformations of the plank. They occur due to differential contraction of the wood within a plank during the curing process.

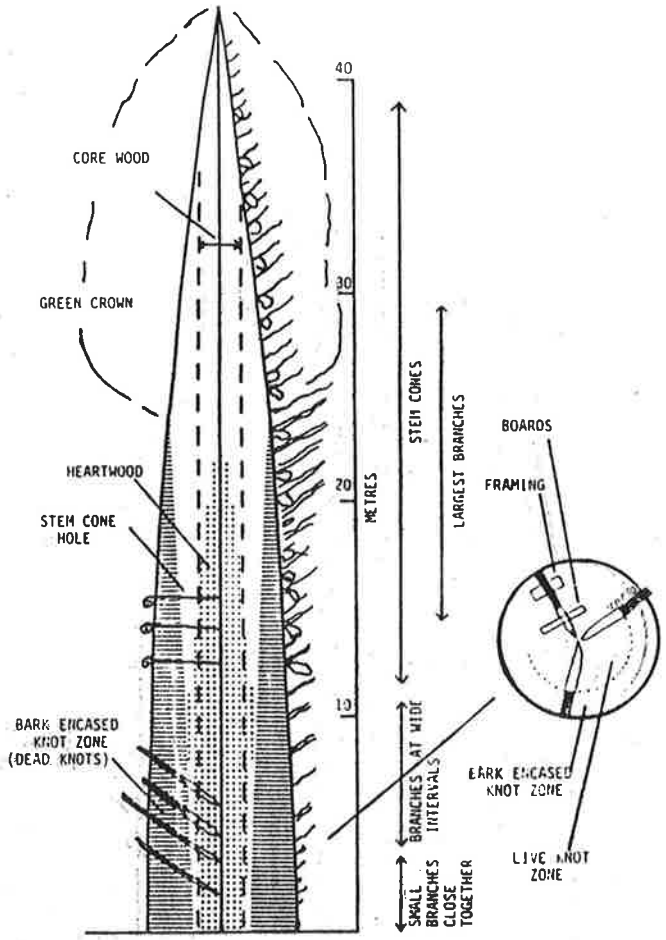

Figure 2.4 Schematic representation of untended radiata pine. (From J.C.F. Walker (1984))

The description of the various features found in radiata pine also requires an understanding of how they are formed. The schematic drawing of a mature tree in figure 2.4 shows the areas of the tree where the features occur. Figure 2.5 shows a section of sawn board that has a number of significant features. The **pith** is the small, central soft cone in the structural centre of the log and is usually between 8 mm to 10 mm in diameter through the entire length and is visible as the small dark circle on the right face of figure 2.5. Each year of growth is marked by bands of varying density called growth rings. These are the **earlywood** and **latewood** and appear on the face of the sawn board as

alternating light and dark lines. The colour of the growth rings is due to the variation of the resin content of the cells and also their density, the earlywood cells formed in spring being lighter than the denser latewood cells formed in summer. These cells transport food throughout the tree.

Figure 2.5 A sawn board showing the features of pith, needle trace, encased knot and growth rings.

The first four or five years of growth form the corewood which is characterised by its lower density and the wide spacing of the growth rings. The inner part of the tree is called the heartwood and its annual rings contain no living material. It is identifiable by its slightly pinkish colour and can be seen in figure 2.6 together with pith and more clearly in figure 2.7. The outer living part of the tree is a straw colour and is called the sapwood. Heartwood is usually but not always of lower density but is always of inferior strength. The formation of the juvenile wood results in the micellar or fibril angle being at a large angle to the vertical compared to later wood. This results in a

greater longitudinal shrinkage during drying than sapwood with a consequence that a board with heartwood on one edge and sap wood on the other will experience a large distortion. This is measured as **bow and spring** as illustrated in figure 2.8. Because of its different properties to the rest of the tree, heartwood is covered by different criteria when used for structural members.

Figure 2.6 A section through the centre of a log showing the pith, pink heartwood, straw sapwood and growth rings.

Figure 2.7 The surface of a board showing an intergrown knot, which is split, the pink heartwood and the straw sapwood. A small amout of bluestain can be detected in the upper part of the board.

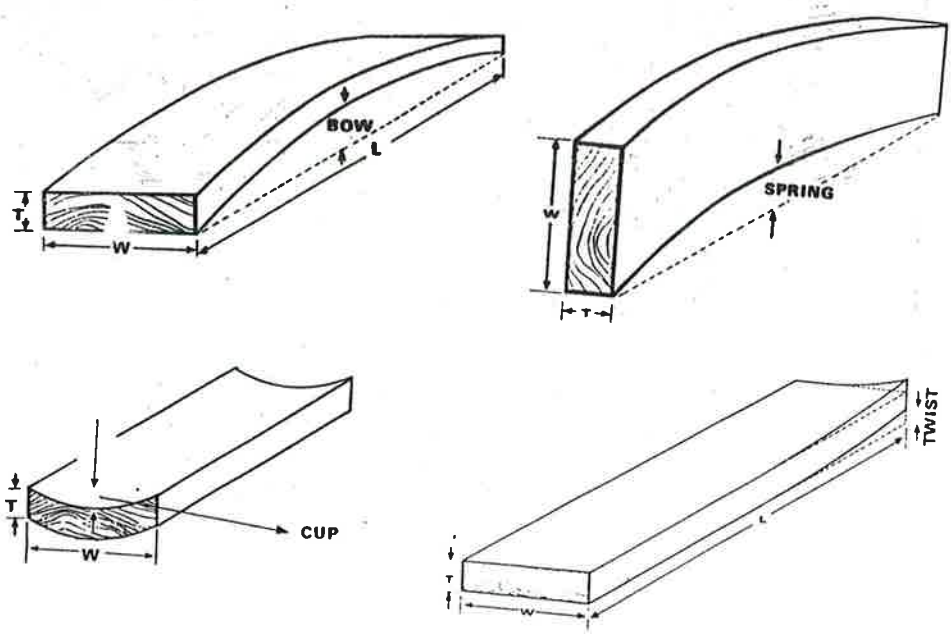

Figure 2.8 Distortion measurements of radiata pine according to grade rule specifications.

The major factor that affects the visual grading of appearance timber is the presence of knots. Branches grow from the trunk of the tree, die, break off and the dead stubs are covered by the growing material of the trunk. The cut branches appear on the sawn boards as oval features called knots which may be intergrown with the material of the board (see figure 2.7) or dead branches encased in their own bark and without structural integrity (figure 2.5). When a branch first starts to form, it has a disrupting influence on the pattern of growth rings. As the branch gets older, it starts to establish its own growth rings (seen on a sawn board as a series of concentric circles rather than parallel lines because of the manner on the cut) and the growth rings of the main body of the tree become more disrupted. The branch also tends to have a higher resin content, averaging 14.5% but varying from 1.9% up to 29.5% at branch death, compared with 2% in the surrounding wood (Reid, 1959). When the branch dies (due to light starvation caused by the growth of the forest) the tree starts to grow around the dead branch (figure 2.5). The branch no longer grows with the tree and the disruption of the growth rings around the branch lessens. However, because the branch is no longer growing there is no bond between the branch and the enveloping tree. This leads to the formation of encased knots in the sawn wood with the branch stub encased by the bark of the branch.

Figure 2.5 clearly shows where a branch has started to grow in the first year of the tree. By the twenty eighth year (from counting the annual growth rings) the branch is dead and has become encased by the growing tree. The lack of a physical bond between an encased branch and the body of the tree leads to a knot in a sawn plank that can fall out which is unacceptable for appearance grade timber. Encased knots are extremely restricted in the grading rules and their detection is a vital requirement of an automated inspection system.

Figure 2.9 Branches becoming encased in the growing tree. Note the bark caught in the angle of the steep branch.

Figure 2.9 shows small dead branches becoming encased by the growing tree. Close inspection of the figure reveals the presence of bark in the angle of the steep branch. This gives a partial encasement that tends to disappear within a few years of growth but leaves a bark pocket that can be unacceptably large. Figure 2.10 shows the effects of defects in downgrading sawn timber in a New Zealand study (Brown (1965)). The effect of bark in the angle of steep branches is recognised as a significant

problem. The process is illustrated clearly in the following figures. Figure 2.11 shows a young knot in the fifth or sixth year of growth of the tree with a large bark pocket. Figure 2.12 is the same branch a few years later. The knot is now fully intergrown. The steepness of the branch can be seen in figure 2.13 along with the annual rings that give an indication of the time scale involved.

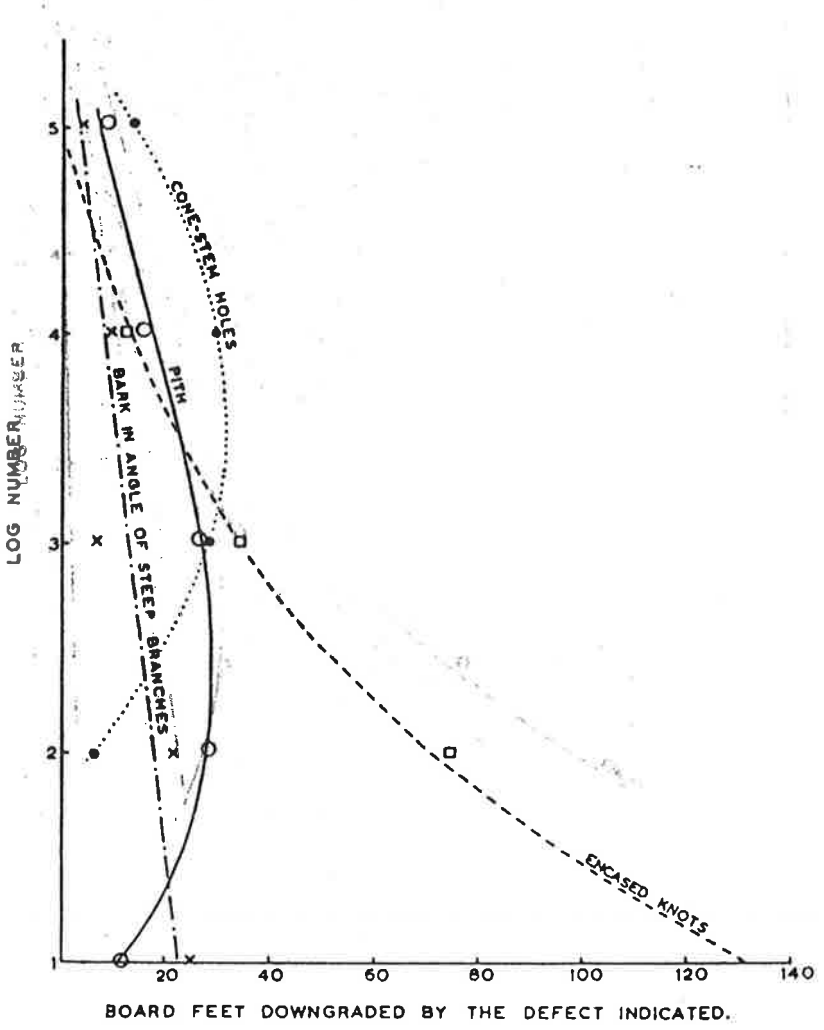

Figure 2.10 Effect of defects in down-grading sawn timber from unpruned Poroporo logs. Total log volumes board feet were: First 222, second 171, third 121, fourth 77, fifth 29.

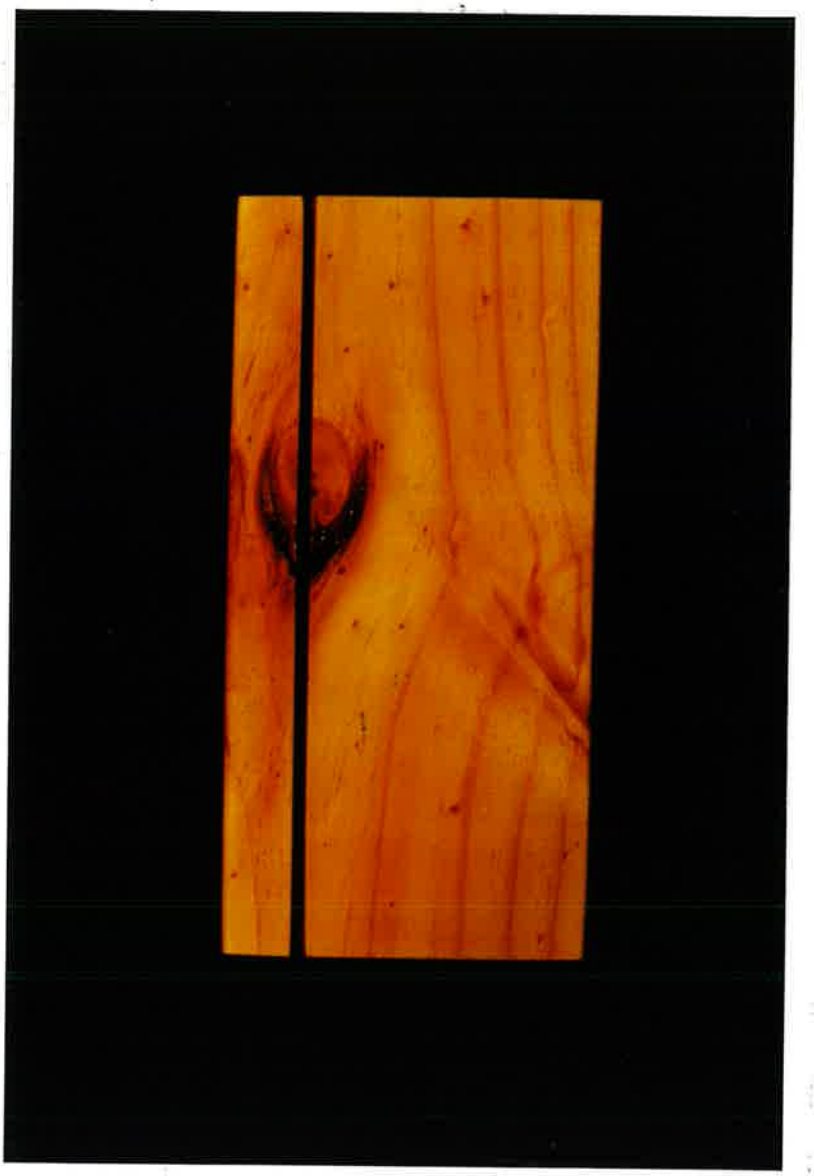

Figure 2.11 A young branch with partial bark encasement. Note the needle trace in the wood which is about six years old.

Figure 2.12 A fully intergrown branch. The needle trace is barely visible.

Figure 2.13 The growth of a branch. Note how the bark has become encased and the branch has become intergrown over a period of four years.

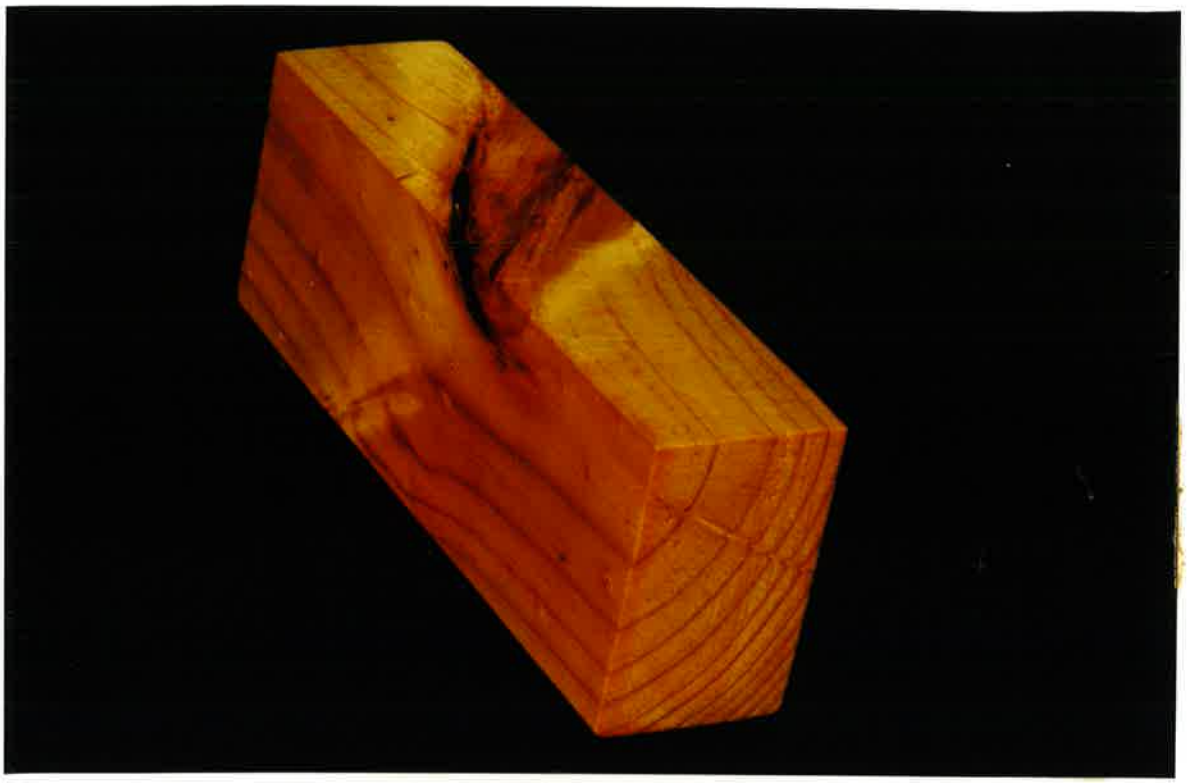

Figure 2.14 The influence of a needle on the trunk of the young tree can be seen in the local distortion of the annual rings that gives rise to needle trace.

Another feature is visible in figure 2.14 as a disturbance to the annual growth rings on the end face of the block. This is called **needle trace** and can be seen on the younger face of figure 2.11 as small dark dots and in the older face of figure 2.12 as barely visible light coloured ovals. The effect is also visible in figure 2.5. It results from the needles on the body of the young tree dying and becoming intergrown in the same manner as small branches. The presence of dead needles can be seen on the bark of a living tree in figure 2.15. The result is pleasing to the eye and does not downgrade the timber as long as it is sound, that is, it is intergrown with the wood and does not have voids associated with it.

Figure 2.15 The dead needles on the trunk of a tree that give rise to needle trace.

Figure 2.16 Pine cones that give rise to cone holes in sawn timber.

Figure 2.16 also shows the pine cones which have a similar effect to that of a dead branch as the tree grows. The point of attachment to the trunk results in small (<10mm diameter) holes in the sawn wood called cone-stem holes. They are referred to in figures 2.4 and 2.10 where their occurrence in the higher reaches of the tree can be noted. Reference to figure 2.3 reveals another feature of the growth of limbs from the tree and that is the regular spacing of branches and their congregation at nodes with relatively clear portions between the nodes. This

effect becomes apparent by inspecting the sawn timber which tends to have features clustered together separated by areas that are largely clear. The exception to this is the base of the tree which tends to have a large number of small branches that result in a large number of sound and **encased pin knots**. The branches that give rise to this can be seen in figures 2.16 and 2.17.

Figure 2.17 The effect of pruning the lower branches can be seen clearly in these young trees.

Occasionally a piece of bark will become encased in the growing wood. The bark of the tree is quite thick and heavily grooved in a vertical direction which is the orientation of bark inclusions. These only form a small proportion of features.

Very wide boards can only be cut from near the centre of the log as the largest size of log is limited, unlike the virgin forests of Canada where logs hundreds of years old are still being cut. The widest board that is cut is 290 mm wide and is very likely to have pith showing on one face. Branches grow radially outward from the pith, which is the centre of the log, so a plank that has a surface near the centre will not only be likely to contain pith but will also cut any intergrown or encased branches at angles close to the axis of the branch, giving the appearance of a long **spike knot** on the suface of the plank.

Figure 2.17 shows the difference between a tree that has had its lower limbs pruned and one that has been left untended. The results in the final sawn log thirty years hence are easily envisaged. Many of the branches will fall off as the tree grows: those that do not will produce encased knots. Figure 2.18 shows the stubs of two pruned limbs in the process of becoming engulfed by the growing tree. Obviously the growth rings will be disturbed for a number of years after the limb has been severed. The significance of this will depend very much on the size of the branch and the care exercised in the pruning operation. Further growth after this initial disturbance will result in clear wood.

The dark appearance of knots from branches that are more than about three years old can be seen in figures 2.7 and 2.11. Figure 2.19 shows a couple of young knots that are not so dark. The middle knot is of a brightness level comparable to the growth rings except for the pith of the knot and it is not easy for a human observer to be definite about the exact outer limit of the feature. Difficulty is also experienced in trying to determine, from the photograph, whether the two lines to the right of the knots and slightly below centre are resin streaks or splits. The middle board of figure 2.20 shows a very faint knot that is over

Figure 2.18 The effect of pruned limbs are evident for a number of years but result in a greater proportion of clear timber.

powered by the dark bark that partially encases it. This figure also shows the difference between wood cut from the straw coloured sapwood (bottom) and the pink coloured heartwood (centre). The features are typical for these two areas, the sapwood containing an encased knot (which has fallen out to leave a hole) and the heartwood containing faint young knots with some encased bark. The upper board is of sapwood and is discoloured by a fungal growth that is called **bluestain**. The range in brightness

over the three pieces is quite large. It is hard for a human observer to discriminate the features of faint knots and splits with diffuse direct lighting and it is no surprise to discover that an automated system also experiences difficulty in this area.

Figure 2.19 Young faint knots. The outer limit of the knots is not easy to define.

Figure 2.20 Top: sapwood with bluestain fungus and knot hole; Middle: heartwood with very faint knot; Bottom: sapwood with large hole, split and encased pin knot.

2.4 APPEARANCE GRADE VISUAL INSPECTION

Grading is performed to maintain a standard between mills manufacturing the same product regardless of the natural variations between different tree stocks. This permits the buyer and seller to be confident of the standard of timber under consideration. There are two distinct grade catagories for radiata pine. Structural grade timber for house frames and other building uses is predominantly mechanically stress graded by machines that measure the force required to induce a given degree of deflection. Although an Australian Standard (AS 1490) does exist for the visual stress grading of structural timber this has been superseded by the increased use of mechanical stress grading.

Appearance grades are established principally on the visual appearance of the wood characteristics and the uses to

which the material is to be put. The grades are based on the premise that only one side will be visible and that this will be the better side. The uses for this timber are largely decorative or require a decorative component in their function and includes flooring, shelving, lining, fascia, joinery and mouldings. These products are covered by a number of Australian standards (Table 2.3) that place them into a number of grades.

AUSTRALIAN STANDARD SPECIFICATION FOR RADIATA PINE PRODUCTS	
AS 1489	SAWN BOARDS (clear, joinery, select, standard)
AS 1492	FLOORING (standard, merchantable)
AS 1493	SHELVING (standard)
AS 1494	LINING and PANELLING (select, standard)
AS 1495	CLADDING, or WEATHERBOARD (select, standard)
AS 1496	FASCIAS and BARGEBOARDS (select)
AS 1497	JOINERY (joinery)
AS 1498	MOULDINGS (mouldings)

Table 2.3 The Australian Standard Specifications that apply to appearance grade radiata pine and the grades of each type.

Sawn board grade timber is not surfaced to a smooth finish and still shows saw marks in the wood for its overall length on the four longitudinal faces. The dimensions of sawn timber are approximately 4mm in width and 3mm in thickness oversize to allow for dressing or profiling to a final size. The other grades are derived from sawn board grades by dressing the timber to a final size and surface finish. Clear timber is graded on both faces and both edges, joinery grade on one face and both edges, and select and standard on one face only. The grade descriptions from the Australian Standard for Sawn Boards (AS

1489) are listed in Appendix 2 and give a good example of the rules that are applied in the other standards.

Clear grade is to be free of all defects on both faces and both edges. The only other aspect to be measured is the amount of distortion in the timber in terms of bow and spring. This is the highest grade of sawn board available for manufacture into components that require a fine knot free appearance and so it is used extensively for mouldings.

Joinery grade timber is reasonably clear timber that is very straight that is used for door frames, door jamb linings, window frames, window sills, etc. The maximum size of a knot is limited to one half the face width and only three knots are allowed in any 4m length. **Select grade** is only slightly less restrictive than joinery grade and allows limited pith. This grade is used for panelling, shelving and other uses where a knotty timber is required that is strong and of good appearance. **Standard grade** is less restrictive again and is suitable for less exacting requirements than that of joinery or select. Timber that is unsuitable for standard grade is classed as **merchantable grade** timber.

The standards specify the minimum quality for each grade so that the average quality in any grade is better than the minimum described and will contain representatives of the spread of quality in that grade. Many pieces will only be restricted from a higher grade by a single grade controlling imperfection.

These, briefly, are the rules and standards for radiata pine that are applied by the graders in the timber mill. The performance of the graders is examined in the next section.

2.5 QUALITY ASSURANCE ASSESSMENT OF MANUAL GRADING

The motivation for this study is to improve the efficiency of the grading process. In order to understand the nature and scale of the problem an assessment of the performance of the manual grading process is performed using the results of the local quality control inspections.

After the timber has been milled to final shape and finish and sorted into grades it is stacked onto a pallette, wrapped in plastic and tied. Quality control inspection involves untying and unwrapping the pack and regrading each piece of timber in the pack in a more rigorous manner than the graders on the line are able due to time constraints. The reports of 25 pack checks are analysed and the results displayed below in tables 2.4 - 2.6.

GRADE	No. of PIECES	PIECES OUT OF GRADE No.	%
MERCHANTABLE	913	219	24.0 above grade
STANDARD	1962	115	5.9 below grade
FURNITURE	380	68	17.9 below grade

Table 2.4 Quality control inspection of 25 graded packs.

Table 2.4 illustrates the broad essence of the quality control problem. Almost one quarter of the merchantable grade timber is actually of the higher and more valuable standard grade. This is a direct loss of revenue as the standard grade sells for approximately 1.8 times the merchantable grade per cubic metre. The standard grade approaches the desired limit of 95% ingrade material from the direction that is most economically beneficial to the manufacturer. Reducing the amount of out of grade material from the standard and higher grades is desirable for consistency but results in a lower return for the product. In particular, because the standard grade is produced at approximately three times the volume of merchantable grade and is worth 1.8 times the price a drop of 1% point in the amount of below grade material in standard grade needs to be balanced by a 5.5% point decrease in the amount of above grade material in merchantable grade to achieve the same return. This assumes that

the prices charged for the better graded material remain constant.

MERCHANTABLE GRADE TIMBER		
PIECES OUT OF GRADE		REASON
No.	%	
73	8	STANDARD GRADE
146	16	RE-DOCKABLE

Table 2.5 Reason for being classed out of grade in 913 pieces of merchantable grade timber in seven packs.

The reason for such a large amount of standard grade timber in the merchantable grade packs can be discerned from table 2.5. One third of the wrongly graded timber can be placed straight into standard grade. This portion of the timber is the result of the tendency for graders to downgrade a doubtful piece rather than upgrade and is evidence of the inability of manual grading to apply the grading rules accurately at the rate of production. The same criticism can also be directed at the other two thirds of wrongly graded material which can be upgraded to standard grade by docking out the feature that excludes the piece. This is done if the remaining piece is greater than half the original piece length. This only becomes of economic benefit if the docked length is over 56% of the original length as the increase in value to standard grade is 1.8 times. The docked material becomes scrap as far as board material and is either burnt for process heat or pulped.

REASON FOR DOWN GRADING	STANDARD GRADE No.	%	FURNITURE GRADE No.	%
PITH	8	7.0	3	4.4
CONE AND KNOT HOLES	46	40.0	18	26.5
WANE	1	0.9	1	1.5
RESIN AND BARK	0	0.0	3	4.4
MACHINING	21	18.3	11	16.2
DISTORTION	0	0.0	6	8.8
VOIDS	1	0.9	3	4.4
SPLITS, CHECKS, CRACKS	36	31.3	21	30.9
BROKEN, DAMAGED	2	1.7	0	0.0
ENCASED KNOTS	0	0.0	2	2.9
TOTAL	115	100.0	68	100.0

Table 2.6 Reason for the down grading of timber that is judged to be wrongly graded during quality control pack check inspections of 18 packs.

The type of features that are inaccurately judged by the graders and lead to timber being placed into a higher grade are listed in table 2.6. Standard grade suffers from the inclusion of a large number of cone and knot holes. Both grades suffer from a large amount of splits and cracks followed by a lesser amount of machining imperfections. These include skip, where the board is not thick enough to allow the milling cutter to provide a smooth finish, and planer burn, where the cutter has resided on a stationary board leading to a black, charred band. Clearly, an automatic inspection system must be particularly capable of detecting these features as they form the majority of downgrading defects.

3. LITERATURE SURVEY

3.0 INTRODUCTION

This survey examines the application of automated visual inspection to the grading of timber. It begins by examining previous and current methods of automating this inspection process. A distinction is drawn between inspection for cutting strategies and inspection for appearance grading. Applications to other industrial situations is then explored paying attention to similarities with the analysis of wood. This is done in conjunction with an investigation into the algorithms used to interpret visual images. The survey concludes with an assessment of the strategies employed in the interpretation of vision information and the hardware implementations available.

3.1 AUTOMATED VISUAL INSPECTION OF TIMBER

The effectiveness of the inspection process is measured by the losses incurred due to incorrect grading. With the inspection of radiata pine timber for appearance grades the consequence of downgrading high quality timber to a lower grade during the examination procedure is a measurable loss in financial return due to the wood being sold for a lower price. The loss associated with upgrading poor quality timber into a higher grade is an intangible degree of customer dissatisfaction. Because the wood is used only for appearance functions and not for structural purposes there is effectively no risk of personal injury liability resulting from misclassified timber. The purpose of the automated visual inspection of timber is to implement an improvement in quality control. This is necessary to maximise the return on the product and to ensure customer confidence in the grading classifications.

The purpose of the inspection determines the method that is used to scan the timber. For example, determination of the optimum cutting strategy requires inspecting the timber before it is processed and then directing the cutting process. The features that are sought are those that degrade the value of the finished timber and are not restricted to surface features. Internal aspects of features also require detection. This has led to research using ultrasound, microwave, X-ray, and neutron, as well as optical methods. A summary of these methods is given in Szymani (1981). Apart from optical devices these transducers respond to changes in density which is directly related to wood strength. Knots appear as concentrated regions of high density, heartwood as a large region of low density and cracks as thin regions of low density. None of these methods resulted in a successful application, and research has been predominantly directed toward optical systems.

Appearance timber is graded on the surface image after the wood has been sawn and planed to a smooth surface. The internal structure of the plank is assumed from this image and so an optical system would seem the ideal choice. Decisions for sawing stategies that are conducted manually at present are also based on the surface appearance and so it is not surprising that optical systems have been developed for this purpose as well. Optical systems can be divided into two catagories: laser scanners and camera systems.

3.1.1 LASER SCANNING SYSTEMS

The first optical systems developed were based on laser scanners. Mueller (1976) and King (1978) describe a system developed at Bendix Research Laboratories based on a He-Ne laser that uses a rotating mirror and several fixed mirrors to illuminate sequentially both sides of a plank with a spot approximately 1mm in diameter. Four photomultiplier tube assemblies amplify the scattered light signal which is processed with high and low pass analogue circuits to accentuate cracks and larger defects respectively. These signals are then

digitised and processed before passing to the minicomputer which encodes the locations and sizes of the defects detected. The scan lines are interleaved with the scans on the other side of the board and are spaced 1.27 mm apart on each side.

The minicomputer forms a list of defect types and rectangles which are used to determine optimal cutting lines. This data controls two automatic saws to remove the defects, one fixed and one moveable for crosscut and ripcutting patterns.

The laser scanner proved able to detect splits, checks and red (tight) and loose (black) knots of various sizes but had difficulty detecting lighter coloured features and acceptable markings are detected as defects. In order to compensate for the deficiencies of the scanner a human inspector is employed to enhance light features with a black felt pen and to suppress acceptable flaws with a reflective paint before the wood enters the scanner. This restricts the performance of the scanner to the skill of the inspector and defeats the purpose of the system to automate the inspection process. The system is no longer in service.

A second laser detection and sawing system is referred to by Conners (1983) and described in the patents of Mathews and Beech (1976, 1977). The system is claimed to be able to detect and perform a limited differentiation of defects such as knots, bluestain and rot using a laser scanner, photodetectors and filters. Being developed by Plessey Co. Ltd. and Weyerhaeuser Co. there is a lack of more detailed information.

The problem with laser scanning systems is that the laser light is monochromatic and so is absorbed and reflected by the wood surface in a manner that is different to that of white light. The feature detection that is done by human inspectors is based on the analysis of images formed from the reflection of a panchromatic white light and so it is not unreasonable to expect that analysis of optical images should prove a more fruitful line of research.

3.1.2 CAMERA SYSTEMS

Research using camera systems has developed steadily since the early 1980's with the increased availability of equipment for the interfacing of video cameras and recorders with small computers. These frame grabbing digitising boards enable an image from a video camera to be stored in digital format and displayed on a monitor. Camera input can be from a linear array or a rectangular array camera.

The most extensive report on research into automated visual inspection of timber is a series of papers by Lin (1983) and Conners et al. (1983, 1984). Described is a method to find and classify defect areas in images of wood. The purpose of the study is to recognise defect areas in order to describe a cutting plan that will yield the greatest quantity of defect-free, clear wood. This wood is to be used for furniture manufacture and the emerging technology of laser cutters will allow a more flexible approach to cutting. The laser cutter can be made to cut any desired contour and so the cut of a plank can be optimised if the location of defects is known.

The species used is Southern Red Oak and while many of the features are similar to radiata pine there are important differences between the species and grading processes which lead to different approaches to a solution. The purpose of the automated inspection of Conners et al is to plan an optimal cutting strategy. This requires knowledge not only of the presence of a defect but of the type of defect present.

The reason for this is the variety of uses to which timber may be put. Visible exposed parts such as a table top require a defect free finish whereas an internal part such as a sofa frame may only require to be free of defects that reduce the strength of the piece. Radiata pine is a very knotty timber and appearance grades are not required to be defect free but of a known standard of defect distribution. In fact the local industry terminology refers to features rather than defects.

The resolution of the defect extent does not appear to be critical for the sawing strategy of Southern oak and this is

reflected in the classification techniques which give a resolution of 25mm. Conners (1983) presents pictures of some common defect classes and of clear wood. Oak does not appear to have the prominent background of growth rings that is the main feature of radiata pine. These growth rings vary considerably in spacing and contrast which complicates the task of detecting and classifying features. The grading of radiata pine requires not only the type of feature present to be identified, but also its size to a resolution of 2mm. This factor imposes the requirement for a different approach to feature recognition than that pursued by Conners.

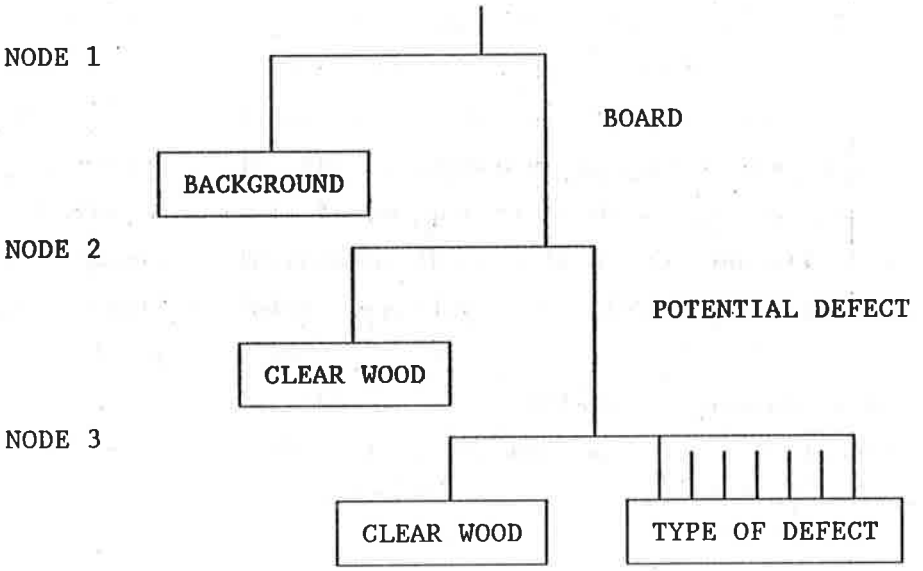

Figure 3.1 The sequential classification strategy used by Conners (1984). Node 1 separates regions containing only board fron those comprised all or in part of background. Node 2 separates clear wood from potential defects. Node 3 attempts to classify the type of defect present.

Conners method is summarised in figure 3.1. The image is divided into 64 by 64 pixel squares representing a 25mm square of wood. Node 1 differentiates the board from the background by assuming the boards lie on a black background. A region is then considered to be background if more than a specified number of pixels are below a threshold value. Node 2 differentiates clear

wood from defects. Conners (1983) shows that the computationally inexpensive tonal measures of mean, variance and skewness are sufficient to perform this differentiation. This is done because 80% or more of the wood is defect free and it is necessary to reduce the load on node 3 which requires the more computationally expensive texture measures. A Chi-squared test is used, assuming a normal distribution, on the measurement vector of the region to assess, to a certain level of confidence, whether the region is clear.

Conners (1983) presents results for two confidence levels which show the relationship between false negatives and false positives. If a high confidence is required that the defects (positives) will be correctly classified, this is the same as specifying a low false negative rate. A low number of false negatives is achieved at the cost of a large number of false positives. If the confidence of detecting defects is lowered, the number of false positives decreases but the number of false negatives increases. The result is a trade-off situation where the cost of false negatives must be balanced against the cost of false positives.

Node 3 uses the pattern related measures of texture analysis defined by the co-occurrence matrices to classify the region into one of ten defect classes. The decision method used is a pairwise multi-class classification scheme whereby, of the K possible classes, K(K - 1)/2 class pair decisions are implemented and the result is the class into which the region is placed the most times. The class conditional density functions are assumed normal and the a priori probabilities of the K classes are all equal. Each class pair decision is made using Bayes method.

Two sets of results are presented by Conners. Conners (1983) gives the results using a sample of 1002 regions of which 192 are clear and 810 are defect regions. Conners (1984) gives the results of applying to a new sample of 1326 regions the classification logic derived from the previous study. This second set consists of 977 clear regions and 349 defect regions to

reflect the ratio expected from the timber used. The results show that the node 2 classification into clear regions and defect regions is consistent at 96.4% and 94.6% correct classification respectively for the two studies.

The results of the ten-class problem are not so consistent. The first set of data shows only one class with a classification accuracy below 76% but the second study has four classes below the 68% rate. Part of this is due to the smaller number of feature regions in the second study with only a few samples in some of the classes. Difficulties seem to lie in correctly identifying dark bark which tends to be classed as knot and hole, and with light bark which tends to get classed as wane, stain, splits and checks. Knots and holes are also confused. The conclusion drawn by Conners et al is that while the node 2 classification into clear and defect areas can be adequately performed using the tonal measures of mean, variance and skewness at the rate demanded by a real time system the ten class classification is not acceptable for an industrial situation. They state that further research is being conducted in the area of colour information.

The node 2 classification of regions into clear and feature areas appears to be a useful approach. Radiata pine has a very predominant background of growth rings which may necessitate the use of spatial information in the form of texture to aid the discrimination. The node 3 classification of regions into ten classes does not seem to be useful. It lacks the classification resolution into the ten classes for the data with which it was developed. Further, the method is limited to a spatial resolution that is the size of the region which is 25mm. The grading of radiata pine requires a spatial resolution of better than 2mm in order to define the size of the features. The approach of Conners appears to be useful for feature detection but other methods are required to size and discriminate these detected features.

3.1.3 HOUGH TRANSFORM METHODS FOR FEATURE RECOGNITION

The problem of the spatial resolution of the features of wood is addressed by Poelzleitner (1986) in a paper that expounds the use of the Hough transform (HT). HT methods were originally developed to detect straight lines in noisy images and were later generalised to curves and shapes (Gonzalez and Wintz (1987), Brown (1983), Shapiro(1978)). The HT is a mapping of the image (or a feature space) onto a parameter space which is an accumulator array. Each point in the image "votes" for a point in parameter space. Points in parameter space that collect high values reflect features in the image space depending on the form of the transform.

The approach developed is to divide the image into squares of 4 by 4 pixels and to use the second order statistics of the grey level distribution (co-occurrence matrices) and a linear classification scheme to place each square into one of 16 primitive classes. These primitives are candidate building blocks that are without indication of which type of feature they are a part of. The 16 classes include things such as "left side of object", "centre of object", "north-west corner of object", etc, and form a feature space.

The method then forms a parameter space that is isomorphic with the image space by casting votes according to the feature spread point function g[C, f, x], where C is the feature class, f is the feature vector, and x is the symbol coordinates. This has the effect of tidying up and confirming the feature points and eliminating isolated points. The major difference that is claimed for this method over other HT methods is that it does not look for a parameterised description of borderlines of objects or for a recognition of certain shapes. It qualitatively indicates for each point in the symbolic domain whether it belongs to the interior of an object or not.

In this sense it is certainly a practical improvement over previous HT methods which tend to be limited in their usefulness by the large dimensionality of the parameter space. The reduction of the image space to squares of 4 by 4 pixels and

the reduction of intensity information to 16 feature classes reduces the size of the parameter space significantly. Success is claimed in separating the features from the background in a series of 100 images but quantitative results are not given. The method appears to offer a degree of promise but still lacks the resolution required for grading. The type of wood used in the study is not described. The use of co-occurrence matrices is restrictive as these are expensive computationally. The method requires their calculation for the entire surface of the board which would place a strain on resources in a real time application. It could conceivably be applied to the output of the node 2 of the Conners system, that is to the areas classed as feature areas, but at this stage the presence of the feature is known and it is the type of feature that is sought. The HT method of Poelzleitner does not discriminate features and so would not be useful at this stage either.

3.1.4 DEDICATED HARDWARE FOR TIMBER INSPECTION

Analysis of wood images in a real time industrial situation requires very fast processing speeds. Information is acquired at rates of the order of three megabytes per second and while most of this is clear wood it all has to be processed. This is only possible with the use of hardware that is dedicated to the processing of images which entails the use of a parallel processor.

A dedicated image processing hardware and software system is currently being developed at the Frauhofer Institute of Information and Data Processing in Karlsruhe, West Germany and is reported by Paul (1988). The system is called VISTA (Visual Interpretation System for Technical Applications) and one of the applications to which it is applied is the detection and recognition of features in wood (Tatari (1987a), Tatari (1987b), Paul (1988)).

The VISTA system is designed to be have a versatile structure to enable it to be used for a wide variety of applications. An aspect of this is the ability to acquire inputs

from a number of sources (not just cameras). Two camera inputs are used in the inspection of wood with three lighting arrangements to highlight different aspects of the features. Backlighting is used to detect open defects such as holes and splits. The same camera is used with direct overhead lighting to detect dark defects such as knots, and the second camera is used with side lighting to enhance voids in the surface due to checks, cracks and unevenness. This is an important feature in that it is extremely difficult (often impossible) to infer depth information from a single directly lit image. Backlighting to enhance open defects unambiguously identifies these features and separates them from dark knots and dark resin or bark. Side lighting causes any depth in the surface to cast a shadow and enables the detection of knot checks, cracks, and voids in bark. This reduces the discrimination demand on the analysis of the directly lit image to that of detecting knots, pith, bark and resin and discriminating and sizing the features.

Three filter modules are used to detect features. The first filter module is for detection with fixed and adaptive thresholding. A fixed threshold detects the bright pixels due to the backlighting through open features. Dark regions are detected by adaptive thresholding. Very dark regions indicate unsound knots, bark pockets and shadows caused by uneveness. Dark regions indicate discolourations caused by sound knots, rots or sap stains.

The second filter module detects roughness and chatter marks by looking for a high frequency of occurrence of local grey level extrema which exceed a given difference between adjacent extrema. Stria, generated by little dents in the cutting blade, are detected by the third filter module which computes the average grey level along lines parallel to the stria. Stria are detected if the average values of the neighbouring columns differ too much.

Detected adjacent pixels are clustered into regions and the geometrical and grey scales of the regions are calculated, eg, the bounding rectangle, area, and mean grey level. The defect

regions are classified by applying a hierarchical decision tree based on cuboid classifiers. A feature vector describes the defect region as a point in feature space and different defect classes are represented as cuboid shaped subspaces. A range query determines the subspace in which the feature vector is positioned. The queries are organised in a hierarchy to optimise the result.

Paul (1987) states that VISTA yielded satisfactory results with an inspection of 400 wood samples but more quantitative results are not published. Only the first filter module has been implemented in hardware and a detection of suspicious pixels of one board is claimed in 150 ms. The second two filter modules are implemented in software but a speed of 10.3 m/s for a 20 cm wide board at a resolution of 1 mm is claimed when the modules are implemented in hardware. This is a rate that will enable real time processing if it is realised.

The use by Paul et al of back, direct and side lighting to differentiate different aspects of the features is a dramatic advance in making the feature identification in wood tractable. The samples of wood used by Paul(1988) are of hardwood and so lack the strong growth rings of radiata pine. The adaptive filter routines may not be so successful when applied to the more noisy images. An advantage of the VISTA system is that it is able to apply the filtering routines to all the wood surface. It does not require splitting the image into smaller subimages as done by Conners et al. Detected regions must still be merged and detected pixels must undergo a filtering to reduce the noise level. The identification of stria is not a problem with the grading of radiata but the identification of open defects is crucial to the grading of appearance timbers. Detection of voids in bark pockets and chipping of knots will be enhanced by the use of side lighting, as will the detection of hit and miss.

Tatari (1987b), in a study to aid the selection of texture analysis methods for automated visual inspection tasks, uses as an example the detection of features in samples of softwood that can be clearly seen to have strong growth rings

similar to radiata pine softwood. Local minima are detected along a row of pixels that crosses the growth rings. Each minima corresponds to a growth ring and if the difference between a minima and its neighbouring minimas is too great it is marked as a defect. As this only detects the minima and not the whole defect an adaptive threshold is used to detect dark areas. The two results are logically combined to identify the defect areas. The algorithmic approach of an adaptive threshold presented by Tatari appears to offer promise in the detection of features in softwood.

3.1.5 COMPUTER GRADING

In the event that an automated visual inspection system can correctly locate and identify all the features on the wood, the task remains to grade the wood according to the industry standards. A paper by Joseph (1985) presents an interactive program to grade hardwood to a British standard. The inspection process is conducted manually and the inspector is prompted by the program to respond with aspects of the board such as length, width, position of knots, etc. The purpose of the system is to enable the employment of unskilled people to perform the inspection task and still ensure that the wood is graded correctly. A BASIC program is presented that applies the grading rules. It is a simple matter to alter the program, first to accept the grading rules of the Australian standards for radiata pine; and secondly, to accept input from a feature list generated by an automated visual inspection system rather than from a human inspector in an interactive format. The grading program of Joseph was developed to improve the quality of the grading process. The success of the program is dependent upon the input from the inspector. The success of a grading program with input from an automated visual inspection system is dependent upon the success of the automated system to detect and classify the features in the wood.

3.1.6 SUMMARY OF AUTOMATED WOOD INSPECTION

Initial efforts to automate the inspection of timber used laser scanners to determine saw cutting strategies by detecting the presence of wane. Systems based on this concept are in regular use in Europe and North America. Attempts to use these systems to detect other features have not been commercially successful. Other transducers (ultrasonic, microwave, neutron beam, X-rays, etc.) have similarly been investigated for cutting strategies but none have achieved a large measure of commercial acceptance. None of these systems is known still to be in use. Of the systems about which there is any published information, resolution is limited both spatially and in the discrimination of features.

More recent efforts have concentrated on the use of optical camera systems and a considerable degree of success has been achieved in the detection of features in real time with the use of dedicated hardware and special lighting arrangements. The classification of features still presents problems and, although several European companies employ inhouse optical systems for saw strategies, a system does not yet exist that is sufficiently effective to act as a visual grader. The difficulty is the nature of the wood surface with features exhibiting a wide variation across types. Reliance upon texture measures in small regions to classify features has proved inadequate.

The problem of visual grading does not lie in the application of the grading rules, but in the detection of the features and in the accurate description of the features. Visual grading requires a greater degree of precision both in identifying the type of features and in sizing them than is required for the purpose of saw cutting plans. It is apparent that correct identification and sizing of the features is the area that demands attention if the visual grading process is to be automated.

3.2 INDUSTRIAL APPLICATIONS OF AUTOMATED VISUAL INSPECTION

In the examination of automated visual inspection applications there are many ways of catagorising the references in the field. Literature surveys (Chin (1982), Chin (1988), Wallace (1988)) reveal an explosive increase in the amount of work published since 1982. These surveys, and in books on the subject (Fu (1983)), divisions are made based upon the algorithm used, the level of processing, or the area of application. The approach adopted in this survey is intended to relate to the problem of the grading of radiata pine. For this reason the grouping of methods, applications and algorithms will be into three principal catagories:

(a) Binarisation.

(b) Grey scale techniques.

(c) Classification methods.

3.2.1 BINARISATION TECHNIQUES

In a paper on the design of automated visual inspection systems, Batchelor (1986) identifies a number of important aspects that are of relevance. These following are the initial points of a systems design:

(a) Mechanical handling of the objects to be inspected.

(b) Lighting, optics and image sensing.

(c) Preprocessing to reduce the data rate.

These three initial aspects of design are inter-related and are treated together in this section. Mechanical handling relates to the positioning of the objects of inspection. Sawn pine boards travel singly through the dressing mill in a longitudinal direction and are discharged onto a conveyor belt. It is envisaged that in an automated inspection system for timber the boards will pass in front of the image sensor in a known position to provide a reliable datum point. This will obviate the need for position detection routines in the image processing.

The lighting will need to be designed to enhance the image captured by the image sensor for the detection of the desired features. Structured lighting has already been mentioned

in the previous section for the detection of holes, splits and surface roughness. The optics and camera system need to be specified in conjunction with the lighting to realise an integrated system. The use of backlighting for the detection of through features requires that the back light be of sufficient intensity to appear as a high contrast white space in the camera image to differentiate it from the lighter regions of the wood surface. Similarly, the gain and contrast levels of the camera observing the side lit image must be adjusted to detect the shadows caused by cracks and chipping over the areas of dark knot and bark.

The whole of the image processing conducted for the purposes of grading timber is an exercise in data reduction. The faster that this reduction can be accomplished the more tractible the problem becomes. The difficult part of image analysis is the reduction of the amount of information into a form that retains the elements that are useful for the grading process. If useful information can be obtained from a simple binary image rather than a full grey scale image then the processing can be speeded considerably. The first levels of data reduction to eliminate useless information is called preprocessing.

The pre-processing stage of binarisation refers to the formation of a bit mapped image, each picture element (pixel) consisting of a single bit, which gives an image consisting of only black or white areas. This is a simplified image and the processing is considerably easier than a grey scale image. For a large class of problems, a binary image conveys all the information necessary to form an accurate description of the object. Examples of the analysis of binary images can be found in Peterson (1985) and Klaus (1986). The following examples perform all the necessary analysis calculations using a binary image. They are presented to highlight the range of applications of this approach and the means used to acquire or enhance the collection of the binary image.

In the Octek Fish Monitoring System (Octek (1986)) whole gutted fish are presented to the camera upon a backlit conveyor

belt to give a clean contrasting image. The binary image is then processed to calculate length and to classify it into one of seven types. Similarly, Cullinan (1984) forms a contrasting image of cereal grains on a dark background to be able to form a binary image. Geometrical methods are used to calculate the axes of each grain and so provide an assessment of the proportion that are broken. These two examples illustrate the use of suitable lighting and handling to ensure a binary image from which all the parameters necessary for correct identification of the features can be obtained.

Baird (1983) uses a binary image of an instrument face to determine the position of the needle pointer using a search wave strategy. Dyer (1983) forms a binary image of the needle pointer of instrument gauges by subtracting two successive images with the pointer at different positions and thresholding the result. This leaves a binary image of only the pointer and further processing using the Hough transform is performed to locate the gauge pointer. The binary image is necessary to restrain the dimensionality of the Hough transform which is used in a form to detect the straight lines of the gauge needles.

Hara (1983) uses a binary image of a printed circuit board to locate defects with comparison to a master image. Optical fibres provide side as well as direct lighting to enhance the image of the circuit tracks whose high specular reflectivity leads to poor image resolution with direct lighting alone. The high contrast image is thresholded to provide a binary image. Thibadeau (1986) thresholds an image of printed wiring boards to form a binary image and the high contrast is achieved by making the illumination coaxial with the camera using a side light and a beam splitter. In a feasibility study of IC chip wire bond inspection Pai (1988) performs all inspection algorithms on binary images achieved using a threshold that is set interactively. Pai identifies an adaptive illumination scheme as a crucial factor in improving the input images and so the reliability of the system.

Aircraft engine combustor assemblies need to be inspected to ensure that the holes are of a correct size. A study by Pai (1986) into the automation of this inspection process uses backlighting of through and blind holes to create a contrasting image. Thresholding provides a clean binary image from which the size of the hole is calculated. In order to detect extruding fins in pulley castings Okawa (1984) places the casting on a dark background to form a contrasting image. A threshold is selected interactively to provide a binary image, the contour of which is used to detect fins. Davenel (1988) describes a system developed to detect bruises on golden delicious apples. Diffuse overhead lighting and a dark conveyor belt is combined with a spectral filter that restricts the light received by the camera to a band that corresponds to the maximum contrast between clear and bruised flesh.

In cytology and haemotology the images of the cells are prepared with special care. Zahniser (1983) describes the refinement of specimen slide preparation that is necessary to achieve a single layer of cells with no overlap and the uniform staining of the cell samples in the examination of pap-smears. This was necessary to allow for the development of tractible algorithms and grey scale information constancy respectively. The use of a fixed threshold for the first segmentation to locate the dark nucleus of the cells is possible because of the uniform staining and the high contrast images. Suspected abnormal cells are detected by the presence of a large nuclear area, measured from the binary data of the first threshold, and are scanned at a higher resolution using the grey scale information for further analysis. In this way, binarisation methods are a way of performing fast preprocessing.

The above examples illustrate that handling, lighting and preprocessing are important factors in the formation of a binary image that will provide a reduced data set for analysis. Mechanical handling must present the inspection item in the correct location and orientation to eliminate or reduce positioning algorithms. Lighting must be structured to give

maximum contrast by the use of backlit conveyor belts; light objects on a dark background; coaxial lighting using beam splitters; and with side lighting to ensure the elimination of shadows. The situation where wood boards are on a conveyor belt requires that the wood be directly lit because the features of interest are on the face of the board. Structured lighting to enhance the contrast of features such as knots, pith and bark is not possible because the contrast of the features is inherent in the wood and cannot be improved by manipulation of the background or lighting. The position of the wood can be controlled to provide a datum point which will obviate the need for an algorithm to find the position of the board in the field of view. Placing the inspection station between two conveyors in a clear space, positioned with rollers at either end will ensure consistent location of the plank for focus and lighting requirements.

The use of structured light to deduce height information in a two dimensional image has found uses in a number of configurations. Ward (1979) describes a system used to guide robots to objects lying on a moving conveyor belt. The objects pass under two focussed slits of white light that are projected onto the same line on the belt. A linear array camera directly above the projected line detects the break in the line caused by the presence of an object. In this way a silhouette of the object is built up, the orientation is calculated from the binary image, and the robot guided to the correct position to pick it up. Nakagawa (1984) describes four configurations of slit projectors and cameras to detect height information and the application of one of these to the inspection of solder joints. Ray (1988) uses an overhead camera and grazing incidence lighting to enhance the images of solder bumps on circuit boards. This allows the detection of features that are not sufficiently contrasted with overhead lighting alone. Brunelle (1986) describes the use of a linear array camera to inspect the connections in printed circuit boards. The camera views the board at an angle and the lighting is also at an angle so that

specular reflections are directly imaged. Surface features create shadows which allow the inference of lead presence. The use of a linear camera and scanning the board to build up an image eliminates the perspective effects that would be present with an array camera.

The use of structured light can be applied to the detection of holes and cracks in timber in a number of ways. The light is focussed into a line incident upon the wood at an oblique angle. A line scan detector perpendicular to the wood would detect the signal from the plane of the wood but a crack or hole would appear as an interruption. A scanner of this sort could have a resolution of 0.4mm if a 512 element array was used over a 200mm board width. The logic would be binary and so could be hard wired into an extremely fast setup. Real time operation could easily be achieved. It could also be expected to detect knot chipping and skip. The wood would have to be guided under the detector in a limited plane to maintain the focus of the imaging system. This type of side lighting is described by Paul (1988) with respect to the detection of features in hardwood.

The detection of features in wood is responsive to binarisation methods. The majority of features are darker than the surrounding wood and their distinguishing characteristics are their shape and dimensions. The shape can be used to identify a round knot from an irregular piece of bark or a long section of pith. The dimensions are used to place the feature within the specifications of the grading rules. This information is easily gained from a binary image. The formation of a binary image is not straightforward, however. The contrast between the features and the background is variable and not under the influence of the lighting. The intensity level for optimum binarisation is variable as the nature of the wood changes from log to log. Bluestain lowers the average grey level of the wood significantly and other variations are the result of events such as prolonged under water storage. This requires the use of the histogram information to enable the determination of a binarisation level. Methods of forming a binary image need to be explored.

There is a class of features that are not amenable to simple binarisation techniques. These are the faint knots, particularly the young knots where the heartwood of the developing branch has a similar intensity to the main body of the tree. These knots will not stand out from the background by thresholding and another solution needs to be found for the detection of them.

3.2.2 GREY SCALE TECHNIQUES

A typical digital image is encoded as an intensity map 8 bits deep giving 256 levels of grey. This information can be interpreted in order to find a suitable threshold for the formation of a binary image to simplify the analysis. The probability of occurrence of grey levels (represented by the histogram) can be used to interpret information for segmentation in a global manner. The spatial relationship of pixels with neighbours of various grey levels can be used to interpret textural information. These methods are discussed below.

3.2.2.1 HISTOGRAM METHODS

The grey level histogram represents the probability of occurrence of each level of grey in the image. A uniform grey image would consist only of one column at the one grey level (pixel 127, for example) and the height of the line would be the count of the number of pixels in the image. A grey scale wedge with an even intensity variation from black to white would be a straight line, with the height of the line being the number of pixels at each of the 256 grey levels and the sum of all the columns being the total number of pixels in the image. Measures derived from the grey level histogram include the mean grey level; the standard deviation or spread of the grey levels; the skewness which describes the symmetry of the histogram; and the kurtosis or flatness of the image. These are referred to as tonal or first order measures. The relationship of a pixel to its neighbours is called texture and there are many measures that describe various aspects of texture within an image.

Tatari (1987b) presents a table of correspondence between defect classes and inspection methods that strongly suggests the use of parameters of the grey level histogram i.e. mean, variance, skewness, and kurtosis. Subjectively, a feature in wood is prominent because it is different from the background due to intensity, shape, texture, edges, size, etc. It is usually darker, and the edges sharper, than the background of growth rings. Statistics of the grey level histogram (tonal measures) would be appropriate to quantify some of these subjective measures.

The use of the grey level histogram to define a grey level for binarisation is called thresholding. A number of threshold techniques are described in surveys by Weszka (1978) and Sahoo (1988). Threshold techniques are only useful if the object and the background have differing average grey levels with a separable distribution. This is the case with dark knots, bark and pith but not with faint knots and bluestain.

A threshold can be applied globally to an image or locally to subimages. Global thresholding suffers from errors if the lighting is anything but uniform across the entire image or if the reflectivity of the subject varies significantly across the field of view. The latter is common in radiata timber, frequently due to bluestain but more often because of the variation in spacing of the growth rings and the density of the resin in the fibres. The histogram of a large image also has a tendency to mask the influence of relatively small features, further limiting the usefulness of the global threshold. Smaller subimages need to be used to enhance the contribution of the features. For this reason thresholding is usually performed locally and the edges of each subimage smoothed if necessary to provide continuous registration. Dividing the image into smaller local areas that are of the order of the feature size enhances the ability of the thresholding to respond to small features.

The threshold can be evaluated from the grey levels of all the pixels in the subimage in a point dependent method, such as the shape of the histogram, or from a local property in the

region of each pixel in a region dependent method. Point dependent methods rely on features of the grey level histogram such as the valley in a bimodal example. If the histogram is not so clear methods based on the energy, entropy or moment of the histogram are used. Region dependent methods rely on texture measures of each pixel in the image to formulate a threshold, the most common of these being the co-occurrence matrices. A comparison of these and other methods are found in the surveys of Weszka and Sahoo who both provide a measure for determining the suitability of the threshold technique. Sahoo defines a uniformity measure and a shape measure to quantify the suitability of a threshold technique and applies it to techniques in the survey in order to compare them using three standard images. These images are of a building, a cameraman, and a portrait of a person. Unfortunately, these are not images that are being processed for a quantifiable end and the threshold suitability is entirely subjective, the uniformity and shape measures having no goal for comparison.

Weszka defines an error measure and a busyness measure to evaluate the goodness of a threshold in a number of Forward Looking Infra Red (FLIR) images which tend to be very blurred blobs of varying overall brightness. These measures are based on region-dependent co-occurrence matrices and are applied to images with a particular goal, that is, to identify a target. Wood images differ from FLIR images in that the texture elements of wood are much larger and the features are not so blurred. Radiata images respond well to thresholding (with the exceptions described previously) using point dependent methods. Region-dependent methods incur the cost of greater computational complexity with no examples to show that it is justified.

The difficulty of finding a suitable threshold measure is mainly due to the range of images that exist for various applications. Some images have strongly bimodal histograms that enable quick accurate thresholding to be performed. Others are unimodal and no threshold is immediately obvious. Only if some knowledge of the image exists a priori can a thresholding

technique be applied. The best technique is application dependent and is difficult to assess from results of unrelated images. For this reason many of the thresholding techniques found in the literature (Dinstein (1984), Sahoo (1988), Weszka (1978), Kapur (1985), Boukharouba (1985), Tsai (1985)) are not found to be directly applicable to the thresholding of wood images.

3.2.2.2 TEXTURE METHODS

Texture measures quantify the relationship of pixels and grey levels to neighbouring pixels and grey levels. Surveys of the different methods available (Haralick(1983), Van Gool (1985)) reveal that of these second order grey level statistical methods, the most popular is the Spatial Grey Level Dependence Matrix (SGLDM), or co-occurrence matrix. These can be used to calculate the measures of energy, entropy, correlation, local homogeneity, and inertia.

Texture measures have been used to characterise an image as a member of set of classes. Haralick (1973a, 1973b) uses the co-occurrence matrices to classify three types of images: photomicrographs of five kinds of sandstone; aerial photographs of eight land use catagories; and multispectral satellite images of seven land use catagories. The results for the sandstone images (89% correctly classified), aerial photographs (82% correctly classified), and satellite images (83% correctly classified) indicate that the texture measures contain essential features for the differentiation of the different classes.

Weszka (1976) compares different texture methods on aerial photographs and satellite images. The texture methods compared are the Fourier power spectrum (PSM); second order grey level statistics (co-occurrence matrices, SGLDM); grey level difference statistics (GLDM); and grey level run length statistics (GLRLM). From a small data base, the conclusion drawn is that the co-occurrence matrices contain a more complete measure of textual features and led to a better classification than the power Fourier spectrum.

Conners (1980) attempts to clarify the results of Weszka (1976) by repeating the comparison of the various texture methods but they are applied to artificial Markov generated textures. While this is a more comprehensive comparison than Weszkas', it is removed from the application of texture analysis to real world images for a functional purpose. However, similar conclusions are reached to the previous study. The SGLDM and GLDM are more powerful than the PSM. The GLRLM is considered poor, while in contrast to Weszka the SGLDM is considered more powerful than the GLDM.

The use of texture measures to quantify the assessment of carpet wear is examined in a paper by Siew (1988). The study compares the methods of the SGLDM, the GLDM, the GLRLM, and the Neighbouring Grey Level Dependence Statistics (NGLDM) and finds most of the measures able to discern differences in wear in the different types of carpets used. Results were only unsatisfactory when, for example, a measure that was suitable for detecting fine texture was applied to a sample with a coarse texture. The advantage of carpet inspection is that the carpet has a regular texture which allows the effect of moderate changes to be detected.

Conners (1978) describes an application of a technique to reduce the size of the co-occurrence matrices. This is done by reducing the number of grey levels in the image. Equal Probability Quantising (EPQ) places an equal number of pixels in each of the reduced number of grey levels. This is shown to be an optimal way of reducing the number of grey levels with the minimum loss of information. The technique is applied to the comparison of radiographic images that have widely differing contrast levels that are to be compared for the presence of a particular texture. The EPQ ensures that the images to be compared have similar histograms without any degradation of the textual information. This is applied to good effect by Packer (1979) in the detection of corrosion in rivetted aluminium panels using radiographic imaging methods.

A methodology for the selection and parameterisation of texture analysis methods is presented by Tatari (1987b). Also described is the application of texture analysis methods to softwood. Softwood texture is described as having a stochastic texel (texture element) distribution because the distances between the growth rings vary irregularly. Because the growth rings vary in grey value and size, and because every growth ring has an inhomogeneous grey value composition the texture is classified as having varying, inhomogeneous texels. The background is classified as homogenous. This describes the growth rings over the background of the lighter wood. If the previous understanding of background is restored (that of the growth rings and the lighter wood together) the description of varying inhomogeneous texels indicates that discriminating features on this background may not be easy as the features, especially the fainter ones, also have varying inhomogeneous texels. Tatari uses the methods of local extrema and adaptive thresholding in the detection of features in softwood which provides a strong incentive to investigate these areas for the inspection of radiata pine.

3.2.3 CLASSIFICATION METHODS

Visual inspection tasks often require the image to be placed into one of a number of classes using the information from tonal and textural measures. Pattern recognition methods allow one to classify an image based on the values of a number of measures. An image is represented by a vector in feature-measure space and a decision boundary places it into a particular class. Pattern classification is covered extensively by Nilsson (1965), Meisel (1972), Patrick (1972), and Bow (1984).

A survey of pattern classification algorithms by Ho (1968) presents concisely the different methods available for different input data. Briefly, there are four types of data knowledge: (a) the functional form of the conditional density function is known for all the classes; (b) the parameters of the conditional density functions are also known; (c) sample patterns

of known classification are available; and (d) samples of unknown classification.

Conners (1983) uses a Chi-square test to make a class-pair decision between subimages containing defects and clear wood using four tonal measures. They then use a Bayes criterion to make a number of class-pair decisions to place a subimage into one of ten classes using the four tonal measures and six texture measures. The Bayes criterion requires that the prior probabilities of the classes as well as the cost of wrong decisions for the error types is known.

The class conditional probabilities are often not known and methods that do not make assumptions of a priori knowledge are finding increasing use with the large amount of data generated from visual images. Blanz (1988) uses and briefly describes a polynomial classifier to classify pixels in images of solder balls with twenty feature measures. This method is fully described in Schuermann (1977).

Bartlett (1988) use a K-nearest neighbour algorithm to classify solder joints into nine classes using twenty seven feature measures. In order to reduce the effect of measures that contributed little to the classification and measures that are highly correlated they use an Objective Dimensionality Reduction (ODR) method to reduce run time computations.

Bartlett (1988) also compares the statistical approach to pattern recognition to an expert system method. This method uses masks to quantify certain aspects of the solder joints, such as central brightness, corner brightness, etc. The results of these masks is compiled giving a weight to each of the different classes. The class with the greatest score is given to the image. The expert system offers several advantages over statistical methods. They can provide an explanation of how they arrived at a particular decision which is helpful for debugging and for modification of the system. They are also more flexible than statistical methods because special rules can be written to handle exceptions.

Miesel (1972) describes a non-parametric classification algorithm that overcomes the limitations of the Bayes method and the K-nearest neighbour algorithm. A linear discriminant function is non-parametric so, unlike Bayes rule, the distribution of the feature measures does not need to be assumed. Unlike the K-nearest neighbour algorithm it incorporates a measure of the distance of a sample from the decision boundary and so is able to move through a period of higher misclassifications to reach a lower level. The K-nearest neighbour algorithm is incapable of this, rendering it susceptible to getting trapped into false minimums. Non-parametric classification requires a training sample of known classification and iterates toward a solution based on a cost function. As long as the training samples are representative of the set as a whole the results will be valid.

Classification methods are only as good as the feature measures that are used to describe the features. Characterisation remains a major open problem (Ho (1968)). Ho also comments that very little work has been done with these algorithms using real data. The work on the classification of visual images in recent years goes a long way to addressing this comment. The use of classification methods in the inspection of timber appears in two levels if the method of Conners (1984) is used. The first is the classification of regions into clear and feature regions at node 2 (Figure 3.1). The second is the classification of feature areas into the different types of defects at node 3. The unsatisfactory results of Conners suggest that alternative methods are required at this stage. The work of Bartlett (1988) with the use of expert system methods suggest that this may hold promise as a way of reaching an understandable solution.

The desire of the Woods and Forests Department in South Australia is to achieve 95% in-grade material which is seen as the level of market acceptability. This requires a combined false negative and a false positive rate of less than 5%. The consequences of false positives (that is, higher grade material) in low grade material is a loss in the realised value of the product, while false negatives, that is, low grade material

finding its way into higher grades, erodes consumer confidence in the product and is undesirable. In the examination of cells in pap-smears the quality control standards are much more exact. Zahniser (1982) describes the development of a system, BioPEPR, that has a false positive rate of 25% and a false negative rate of 3%. The consequences of false negatives in the scanning for cancerous cells can be fatal and this level was considered unacceptable - hence the need for more rigorous standards. The ability to judge the performance of a classification scheme in this way is necessary in order to calculate the potential costs and benefits of the system.

3.3 SUMMARY OF LITERATURE

Initial research into the detection and identification of visual defects in wood was for the purpose of developing sawing strategies, i.e. to set the cutting pattern in order to optimise the value of processed timber (Mueller (1976), King (1978), Lin (1983), Conners (1983), Conners (1984)). This does not require such an exact resolution of feature size and type as is required for appearance grading. These studies have revealed that the detection of features can be achieved to a high degree of reliability but that the resolution of these features is an area that still requires further work. Recent studies into the inspection of hardwood for defects using specialised equipment reveal that processing of the visual information can be achieved in real time (Tatari (1987a), Tatari (1987b), Paul (1988)). These studies also suggest that structured lighting arrangements such as back lighting and side lighting will enhance the detection of various classes of features. Automating the application of the grading rules is a relatively trivial problem (Joseph (1985)).

The survey of literature reveals that schemes to detect features in wood do exist although the exact method for the detection of features in radiata pine requires particular study because of the strong background of growth rings. Codifying grading rules has been performed although not in a manner directly applicable to the Australian standards for radiata pine.

The use of automated visual inspection for the purpose of visually grading timber to a set standard has not appeared in the literature and requires methods different to those developed for the calculation of sawing strategies. The significant obstacle to the realisation of the automated inspection of radiata pine is the accurate discrimination and sizing of the features.

4. APPROACH TO THE AUTOMATED VISUAL INSPECTION OF TIMBER

The aim of this study is to demonstrate the feasibility of automating the visual inspection of pine boards. This is conducted in an off-line environment using a personal computer with a video camera and a frame-grabbing card. It is not intended to develop a real-time system as this would be beyond the capabilities of the equipment. Instead, it is intended to demonstrate that the complex, decision making processes involved in the visual inspection of wood can be codified into a reliable and robust set of algorithms. The possibility of making this process real-time then becomes a hardware implementation of this logic. A possible hardware implementation is proposed.

The approach of the author to the automated visual grading of timber is to divide the process into three steps. These are :

1 - **Feature detection.** The entire area of the plank is inspected for the presence of features. Most of the plank area is clear wood so a quick, computationally cheap method needs to be proven to reduce the data to the feature areas.

2 - **Feature discrimination.** Feature areas are examined closely to determine the type of feature present. The effectiveness of this step is dependent upon the efficiency of the feature detection. It is also the most critical step that determines the effectiveness of the system as a whole.

3 - **Grading.** The plank is graded based on the type, number and distribution of the features according to grading rules. The execution of this step is relatively straightforward but is totally dependent upon the previous two steps.

The line of research that makes up this thesis follows these three steps.

The decision to automate an industrial process needs to based on knowledge of what improvement can be expected, that is, the system must have a known, quantifiable performance specification. There must be some way of judging the level of performance of the automated inspection system that can be compared to the manual inspection. In this study the performance of the automated sytem can be directly compared to the requirements of the grading rules specified by the Radiata Pine Association of Australia (RPAA). In this way there is an absolute measure of performance. An automated system can apply the rules faultlessly if the features are assessed accurately. The performance of the automated system can therefore be calculated from its ability to accurately detect and identify the features that contribute to the grading of the planks.

The effectiveness of the methods used to detect and identify features in pine boards is assessed by the comparison of the known input examples with the output of the analysis program. The input consists of a set of digital images of radiata pine boards. The output consists of a set of simplified images from which the location, size and type of each feature present in each image can be deduced.

It is postulated that the grey level histogram of the image of the wood is sufficient to discriminate the features from the background of prominent growth rings. The image of the wood is digitised from a charge couple device (CCD) camera and segmented into smaller sub-images. Statistics of the distribution of pixel intensities within each sub-image place it in a feature measure space. A decision boundary is derived from a labelled set of samples that divides the feature measure space into two half-spaces; one containing sub-images that enclose a part of a feature, and the other containing sub-images that hold only growth rings. This method allows over 95% of the feature sub-images to be correctly identified. The use of texture measures to discriminate the feature sub-images is found to improve the performance of the decision boundary over that derived using the histogram statistics only marginally at great computaional cost.

Sub-images that are identified as enclosing a part of a feature are merged together and the grey level histogram of the larger feature area is postulated to provide the necessary information to characterise most features. This is supported by the analysis of a series of images containing a range of features which reveals the strengths and weaknesses of the various methods used to discriminate them.

Direct overhead lighting used in this study is unable to discriminate features that contain depth such as splits, cracks, knot chipping, skip and some holes. These features can be detected using side lighting and back lighting and the design of such a system is presented. Also described is the design of a hardware system for the real-time grading of wood in the production environment. This includes the use of ultrasonic sensors to measure the plank deformations of bow, twist, spring and cupping which are whole plank measurements used in the grading process.

5. EQUIPMENT

The aim of this study is not to achieve the real-time processing speed of a prototype system. Instead, the aim is to investigate which image processing techniques can be successfully applied to the automated visual inspection of radiata pine timber. By identifying the low-level processing routines in an off-line environment considerable time and money can be saved when special hardware equipment is built for a prototype.

Figure 5.1 Image processing equipment used in this project including NATIONAL F-10 CCD camera, RAMTEK colour monitor, and IBM-XT clone micro-computer.

The software and image analysis techniques are implemented on a personal computer based system consisting of an IBM-XT clone microcomputer with an 8087 maths co-processor and a

20 Mbyte hard disk drive (figure 5.1). The images are captured with a CCD video camera and stored on VHS video tape, hard disk and in the frame grabbing board memory (figure 5.2). The stored images are displayed on a RAMTEK colour graphic display monitor.

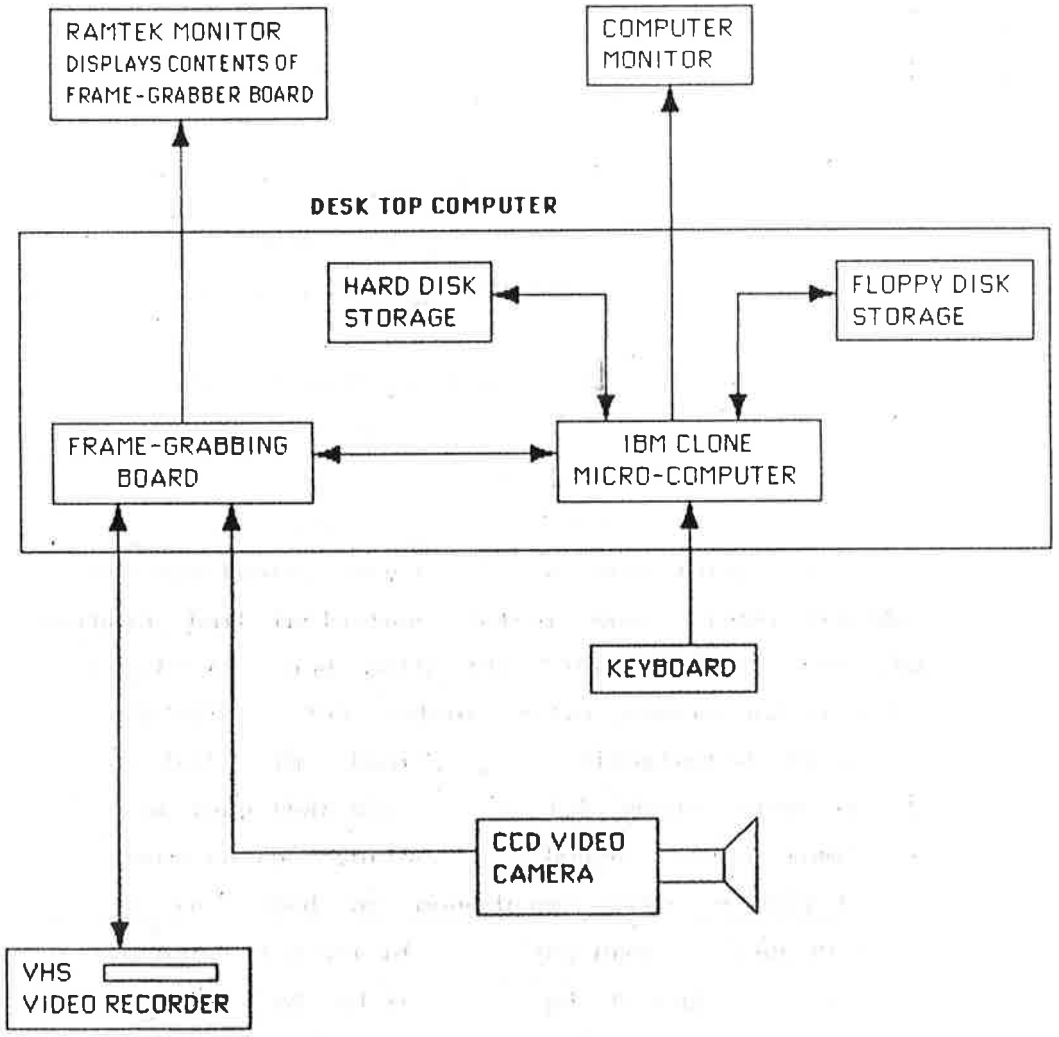

Figure 5.2 Schematic diagram of the image processing system used in this project.

Two frame grabbing boards are used because the first was stolen halfway through the project. A frame grabbing board is simply an array of memory with an analogue to digital input to convert the video signal for storage and a digital to analogue output to be able to display a stored image on a conventional

monitor. Image information passed between the frame grabber and the host computer remains in digital format.

The first frame grabbing board is an IMAGE-ACTION which has a frame memory buffer of size 512 * 512 pixels, at a resolution of 8 bits giving 256 grey levels. The buffer is divided into 256 * 256 quadrants. This was replaced by a MATROX PIP-1024 board with a frame buffer of 1024 * 1024 pixels of resolution eight bits. This is divided into 512 * 512 quadrants of which one can be viewed at a time. Apart from the size of the frame buffer the two boards have essentially identical characteristics. Image data can be passed to and fro between the board and the computer to enable processing of the image information.

Samples of radiata pine boards were obtained from the Mount Gambier timber mill of the South Australian Department of Woods and Forests. These samples were chosen to contain a wide range of features such as light and dark intergrown knots, knot holes, encased knots, very small to very large knots (101 knots altogether); bark inclusions, clear wood, resin streaks, blue-stain, cracks (8); and pith (9). These features were digitised using a IKEGAMI ITC-350 colour video camera and stored on hard and floppy disk. The boards were illuminated by a single 1kW incandescent lamp that was placed 1.8 metres from the board and at 45 degrees to the surface. Although a colour camera is used, the images are stored as monochrome with 256 grey levels. This first data base is a set of 86 images processed on an IMAGEACTION frame grabbing board and is referred to as the FEAT library of images.

A second data base was compiled from a segment of video tape recorded at the Mount Gambier wood mill. A pack that had been processed, graded and wrapped into a bundle of 130 lengths, each of 6 metres, was pulled apart for a quality control check. A number of feature areas covering a range of feature types was presented to a National F-10 CCD camera and stored on VHS video tape. The planks were illuminated by a single 1kW incandescent lamp positioned 1.8 metres from the surface at an angle of 45

degrees. Images were later digitised from the video tape recording to form a library of 48 images. A subset of these images gives a total of 1073 local areas to a training set and 604 local areas to a test set. These images are processed on a MATROX PIP-1024 video digitising board and are referred to as the WOOD library of images. These are presented in Appendix 1A (training set) and Appendix 1B (test set).

The purpose of a training set and a test set is to use the training set to derive parameters for the inspection algorithms and then to test the 'goodness' of these values with the test set. Both sets originate from the same recording session under the same conditions and the division between the two sets is made in such a manner that a nearly equal number of each type of feature is present in each set. Details of the make-up of the two sets is presented in tables 9.1 and 9.2.

A characteristic of the MATROX image store is that when an image is captured in digital format, either direct from the camera or from a video tape, the capture process is not ideal. A black (zero grey level) line, 3 or 4 pixels wide, is formed on the left side of the screen. This is a serious impediment to the analysis of the images and so only a part of the whole 512*512 image is actually stored on hard disk. Software was written which allows one to adjust the size of the part of the image to be stored with the increment in side length in steps of 64 pixels. In this way each stored image in the WOOD library contains only wood and no background area. This can be seen in the images of the WOOD library in the figures.

It was initially hoped to develop the analysis algorithms using a representative sample of wood by recording the output of a production run on video tape and processing this frame-by-frame off-line. The National F-10 CCD camera was chosen because of its ability to capture an image in an interval of 1/1000 of a second (still at a frame rate of 25 frames/sec). This fast electronic shutter speed allows the freezing of motion. The camera was placed over the conveyor belt that transports the surfaced planks from the milling cutter at a speed of 2 m/s. The

camera was adjusted to a field of view of about 280 mm which means that each point on the wood appears in three to four consecutive frames on the video tape. The resolution is determined by the amount the wood moves during the 1/1000 of a second of each frame capture, which works out to be 2mm. Several kilometres of wood was recorded in this way over several hours.

A National AG-6500 editing video recorder was obtained for the play-back of the recorded images. This unit has the facility of displaying the play time and frame number, as well as allowing the tape to be moved forward and backward frame by frame. It was intended to freeze a frame, digitise this frame, and then store part of this frame so that a continuous map of a whole plank would be stored in digital format. The conveyor belt was marked with tape to provide a known length reference to aid in the aligning of consecutive segments of the plank.

When an image frame is stored in analog form on a video tape it is in an interlaced format. The image is stored as two fields, the first contains every second line of the image frame and the second contains the other half of the lines. This format allows a frame refresh rate of 25 frames per second to be displayed as 50 fields per second which the human visual system finds more comfortable. When a video recorder displays a frozen frame it displays a single field 50 times per second. A television monitor has no trouble interpreting this but a MATROX PIP frame-grabbing board loses the sync signal and the picture rolls across the screen. The result is that images can not be digitised from a video recorder in freeze-frame mode but only in play mode. Further, the image that is digitised by the MATROX PIP board contains the second field from one frame and the first field from the following frame. If there is any motion between frames the result is a flickering image. The only way to digitise an image satisfactorily from video tape with this frame-grabbing board is from a stationary image with the video recorder in play mode.

The WOOD library of images were recorded in a stationary postion for about 10 seconds so that they could be digitised with the recorder in play mode.

5.1 CAMERA AND LIGHTING CHARACTERISTICS

Lighting was provided by a single 1KW incandescent lamp in the digitisation sessions for both the FEAT and WOOD libraries of images. The lamp was situated 1.8 metres away from the wood at an angle of 45 degrees to the normal to restrict the purely specular reflections. An aspect of the appeal of radiata pine is the way it catches the light and changes its appearance as the relationship between the source and the observer changes. This is due to the way the light is reflected from the individual wood grains. It was found difficult to eliminate all specular components when the wood is illuminated with a single lamp. At the time that the images were recorded it was felt that the illumination was sufficiently diffuse to provide images of a satisfactory standard. Later analysis suggests that this setup is far from ideal and that results could be improved marginally with more attention to the lighting. The variations in light level across the image as a result of using a single illumination source are found to be of similar magnitude to the variations induced by the camera characteristics.

Both cameras used automatic gain control and so background levels could not be used to compensate for lighting variations across the image as the response of the camera is different for each piece of wood and is different again for any even background. The response of the camera and lighting configuration is investigated in the following series of figures. In the process of capturing and digitising an image from a camera there are two variables of the input signal that can be controlled. These are the gain and offset (contrast and brightness respectively on ones home television set). These variables can best be understood by reference to figure 5.3 which shows an image of a wood board with some pith. The images are obtained using a NATIONAL F-10 CCD camera with the automatic gain

control (AGC) switched off and four different settings for the iris control which determines the amount of light that reaches the sensing element. The object is illuminated by two 1KW incandescent lamps at 45 degrees to the plane of the wood on either side of the camera at a distance of one metre.

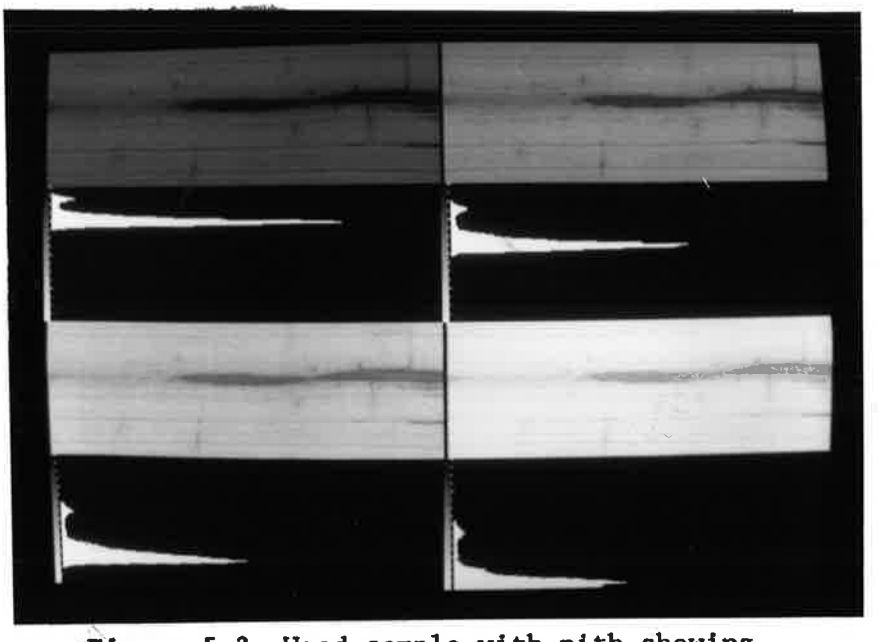

Figure 5.3 Wood sample with pith showing the effect of different iris positions.

Below each image is a histogram where the length of each bar represents the number of pixels (digitsed picture elements) in the image that have the grey level brightness of the corresponding position on the grey scale. The histogram exhibits two peaks, the smaller peak in the darker area of the grey scale corresponds to the dark pith in the image while the larger peak in the lighter area of the grey scale corresponds to the larger area of light background. As the iris is opened more light reaches the CCD sensing element and the histogram is pushed up to the lighter region of the grey scale. The image on the bottom right is so far into the lighter region that the lightest areas of the image are clipped to the maximum brightness level of 255 and information in these areas is lost. This opening of the iris corresponds to the offset (brightness) of the image. It can also

be seen that as the iris is opened and the brightness increases the distance between the peaks increases and the height of the peaks decreases. The histogram becomes more spread out and the information is spread over a larger range of pixels. This is the effect of an increase in gain (contrast) which is performed by the camera with the AGC off and only the iris being altered.

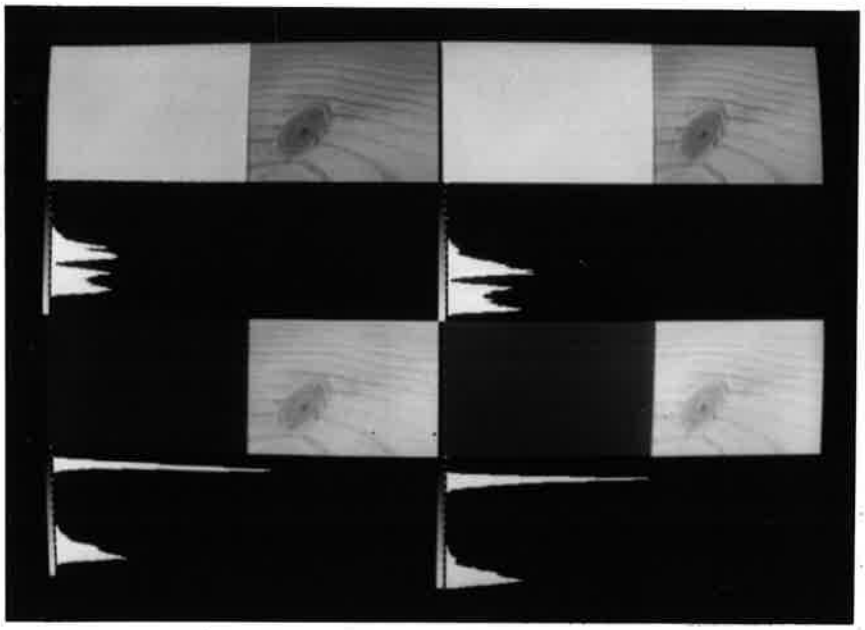

Figure 5.4 The effect of the automatic gain control and contrasting backgrounds.

The influence of the AGC and the effect of image content on the camera response is investigated in figure 5.4. The lighting is identical to that of figure 5.3. The left hand pair of images are with the AGC switched on and the right hand pair are with the AGC switched off. The effect of the AGC is to contain the information in the image within the limits of the grey scale. The bottom right image without the AGC has lost a certain amount of information due to the gain being too high. The upper images contain a piece of white card in left half the image and a piece of wood in the right half of the image. The lower images contain the same piece of wood in the right half of the image but have a black piece of card in the left of the image. In each the peak associated with the respective half of each image

is easily discerned. The position of the peak due to the wood is shifted considerably due to the influence of the peak of the card. The white card shifts the wood image to the lower, darker part of the grey scale while the black card shifts it toward the higher, lighter part of the grey scale. That this occurs whether the AGC is on or off indicates that the gain is being altered within the camera independently of the AGC. With this camera there appears to be no way to capture images independently of the gain control. This fact must be considered in all the processing performed on the images captured using this equipment.

The implications of this to the inspection of timber is that timber of different overall reflectivity will trigger a different response to the automatic gain control and so cannot be compared in terms of absolute grey level values. This places a restriction on the use of fixed threshold levels for different images. A consequence of this is that changing light levels due to bulb failure, brown-outs, or some other physical cause will be automatically compensated for by the gain control. In this way the camera acts like a biological vision system that is tolerant to wide fluctuations in intensity level without degradation of performance. This is the aim of an industrial inspection system - to be tolerant to a changable environment without loss of performance. The use of an automatic gain control on the camera used in this study forces the adoption of certain types of processing which it is considered to be of positive benefit for the directions of the project.

The effect of lighting on the images of radiata pine is demonstrated in figure 5.5. Again the left pair of images are with the AGC on and the right pair with the AGC off. The bottom pair of images are with the two 1 KW incandescent lamps at 45 degrees and a distance of one metre as in figure 5.3. The top pair of images are with only one of the lamps operating (the right one). The fibres in the wood exhibit specular reflections at certain positions which gives a bright area in the right of the image which leads to more spread out double peaked histogram. It is found that these specular reflections are difficult to

control using single lamps as the position of the bright spot will vary with the position of the lamp. The AGC does ensure, however, that the information lies within the limits of the grey scale.

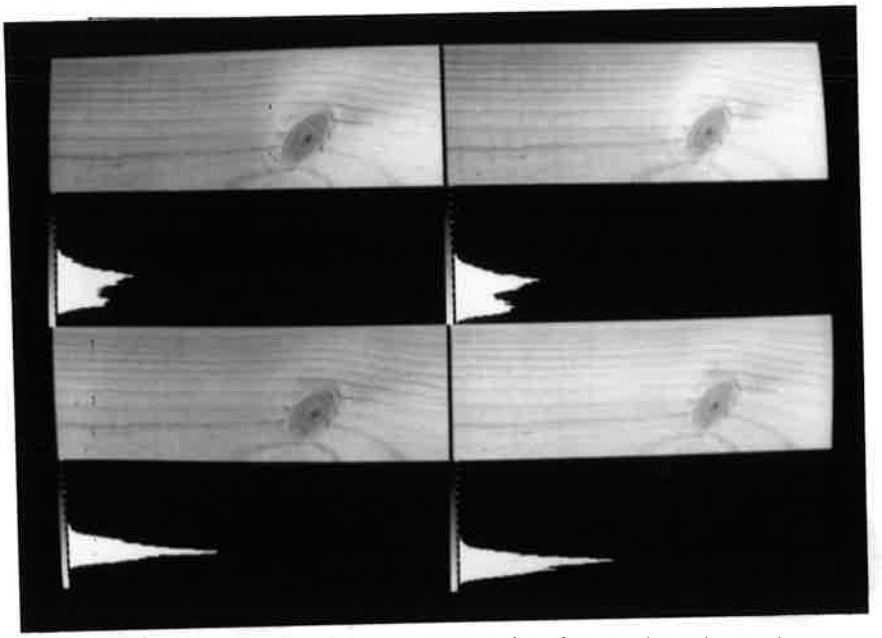

Figure 5.5 Lighting. Single lamp (top) and two lamps (bottom). AGC on (left) and off (right).

A certain amount of specular reflection is visible in the WOOD library of images in Appendix 1. This is despite the lamp being mounted at a greater distance in an attempt to ensure an even lighting intensity across the image. This is detrimental to global image analysis operations but will have a lesser effect on local operations. The AGC was in operation during the collection of the WOOD library of images and all the image information lies within the limits of the grey scale in all the images. Each image of timber filled more than ninety five per cent of the field of view of the camera on a background of a darker table so the influence of background effects as demonstrated in figure 5.4 is minimal and constant.

5.2 ALTERNATIVE LIGHTING ARRANGEMENTS

The use of alternative lighting arrangements to highlight certain features such as through holes and cracks requires a stricter control of the gain and offset controls of the camera response. The priciples can be demonstrated to good effect using the NATIONAL F-10 with the AGC and are presented in the following figures.

Through board features such as holes can be detected by turning down the gain and offset so that no part of the reflected light from the wood reaches the maximum digitised white value of 255. This value will only be reached by light from the lamps behind the wood reaching the camera because of a hole. The following images are taken using a single 1 KW incandescent lamp at an angle of eighty degrees to the plane of the wood and at a distance of one metre.

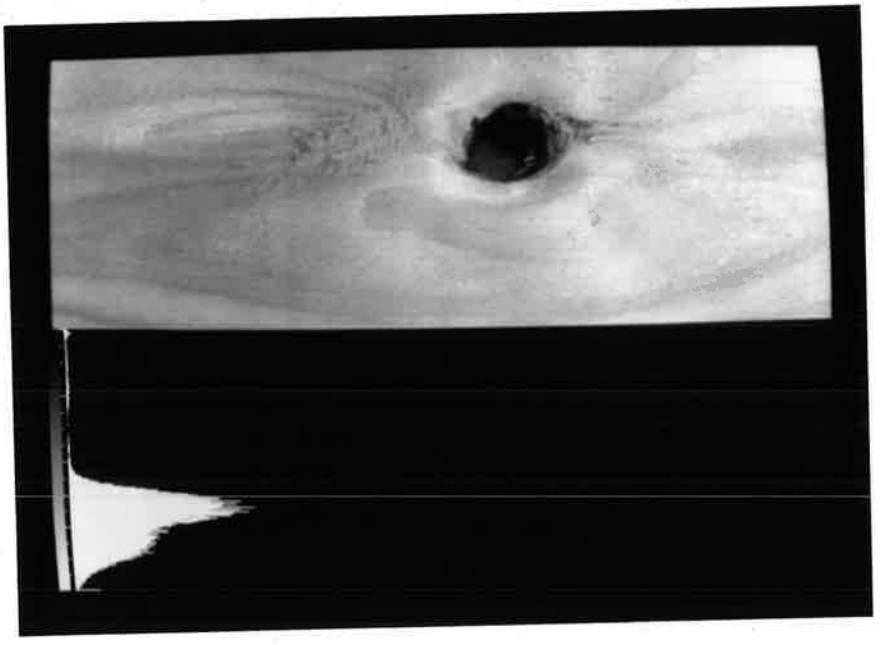

Figure 5.6 A hole in a board with a dark background.

Using a black card background, figure 5.6 shows a knot hole to contain a wide range of dark pixels as revealed by the trail of pixels in the low dark part of the histogram. Ambiguity arises because the shadow caused by the angle of the lamp is darker than the background card that is lit by the lamp while the

bark pixels are spread over a large range. It is important to be able to distinguish the hole from the bark because bark is treated differently to holes in the grading rules.

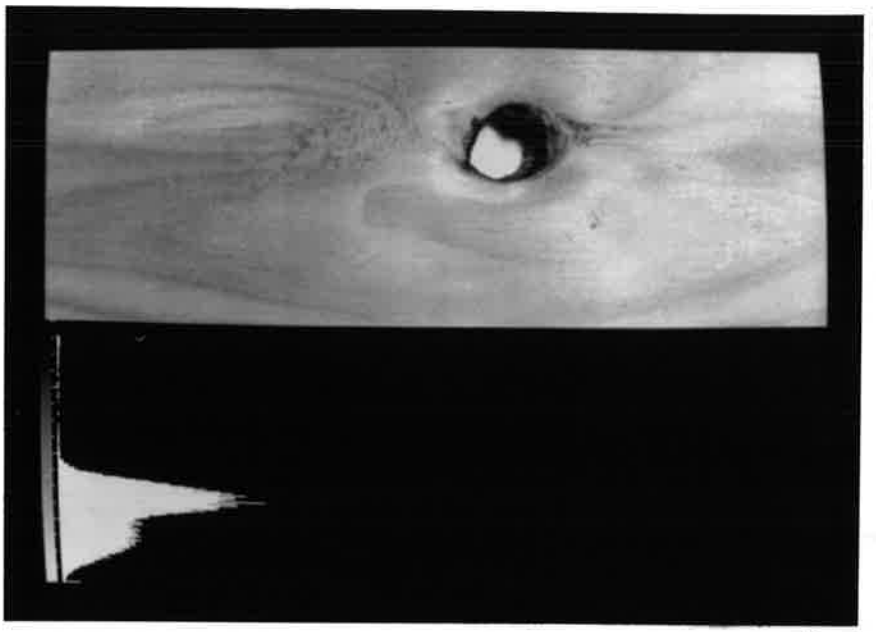

Figure 5.7 A hole in a board with a white background.

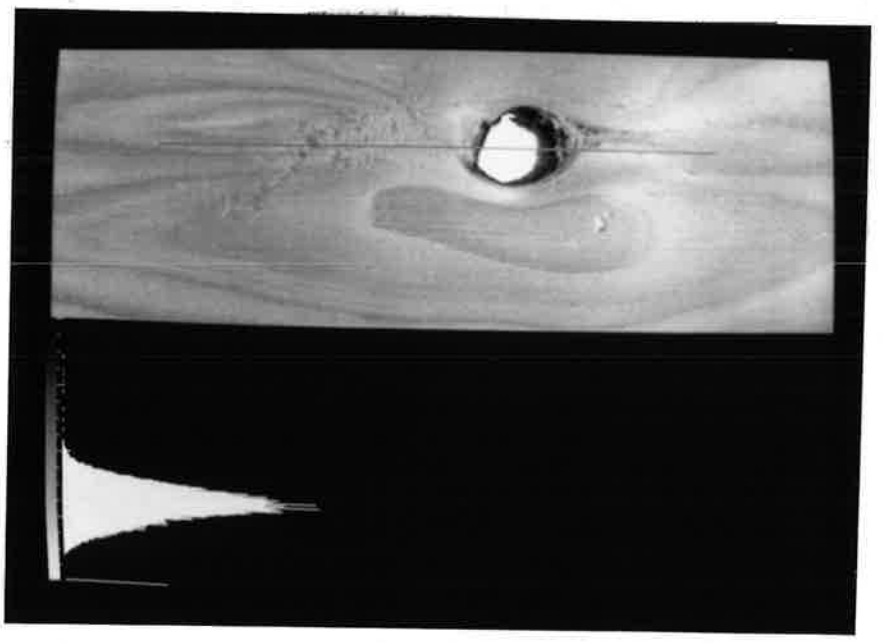

Figure 5.8 A hole in a board with back lighting to highlight the feature.

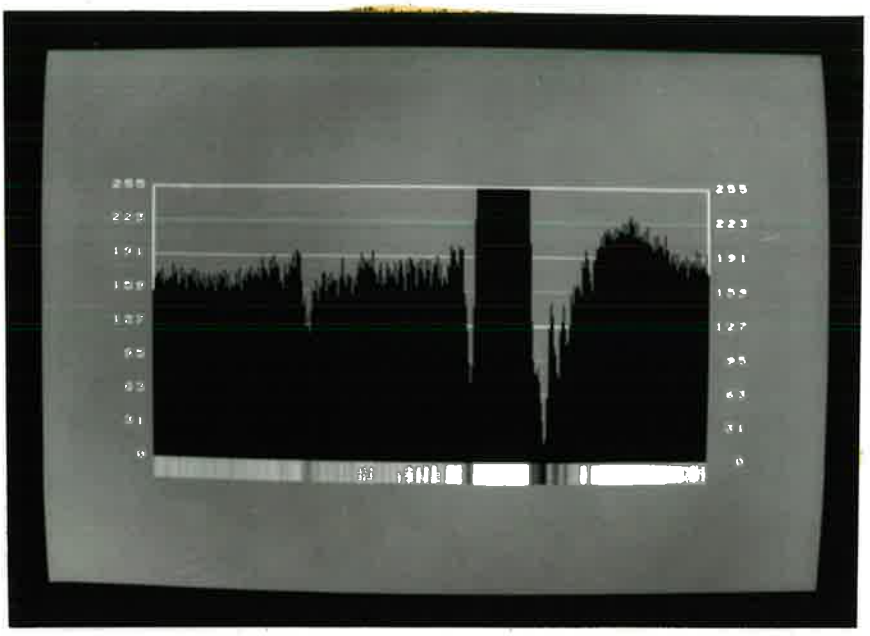

Figure 5.9 Pixel values of the line in figure 5.8.

A white card background in figure 5.7 is little better. The shadow caused by the lamp is dark yet the card in the hole is not easily distinguished from the wood which, from the histogram of this and the previous figure, can be seen to contain pixels of the maximum white value of 255. A clear contrasting result is obtained by backlighting the white card with another lamp as shown in figure 5.8. This has the effect of reducing the gain of the AGC and pushing the histogram toward the lower darker region of the grey scale leaving the hole as a feature of a single grey level intensity of the maximum value. The bark is a seperate feature in the lower darker part of the grey scale. The line across the hole in figure 5.8 marks a row of pixels that are examined in detail in figure 5.9. This shows the grey value for the pixels on the line and reveal the hole to be at a value of 255 all the way across. The bark can be seen to go down to a value of 16 adjacent to the hole.

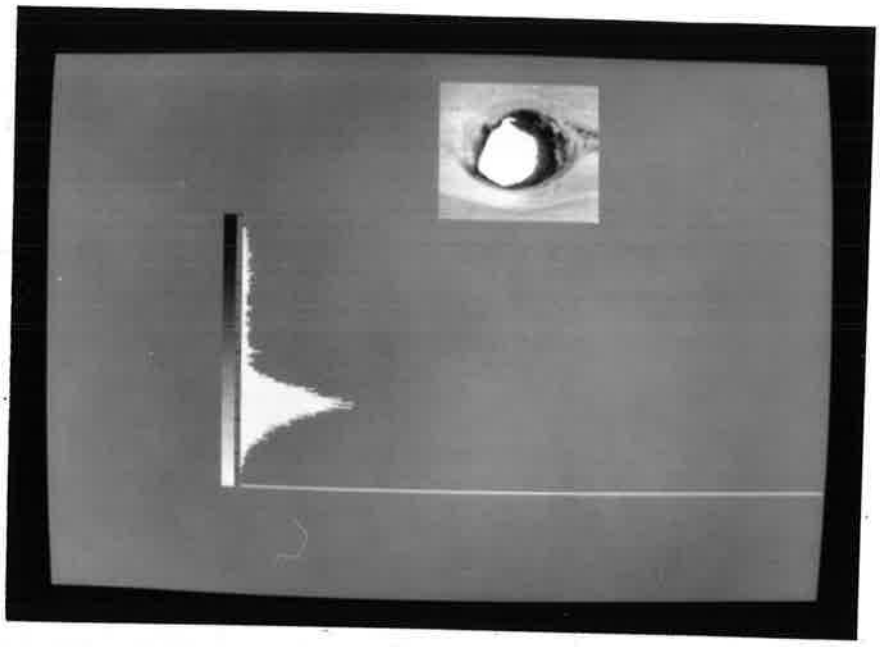

Figure 5.10 Histogram of the backlit hole of figure 5.8.

The histograms in figures 5.6, 5.7 and 5.8 are all to the same reduced scale and do not show values of less than 20. Figure 5.10 shows the histogram of a reduced area which includes the knot to a larger scale. The range of pixels that make up the bark can be seen to extend continuously from grey level 10 to the start of the background peak, a fact which is not obvious from the reduced scales of the previous figures. The distinct break between the board information and the backlit hole is clear and means that the two can be seperated by a threshold, that is, that every pixel of value 255 must be due to the back lighting and so represent a hole. Note that this can only be done if the gain and offset of the camera are set so that at no time does the wood image contain information at the level of 255.

Similarly, when side light is used to highlight cracks and roughness the gain and contrast are turned up so that the dark shadows will appear darker than the dark features of knots and pith and will be the only features to have a digitised value of 0. This type of lighting is very sensitive to stray diffuse light and will need to be enclosed.

5.3 PROCESSING SPEED

This PC based system has the advantage of being flexible and allowing many different ideas to be tested quickly and independently in an off-line environment. The PC is not a viable option for a real time system because of its inadequate capacity and speed. For example, the present processing mill dresses boards at the rate of 2 m/s. If a resolution of 0.5 mm/pixel is sought for a board 240 mm wide then the raw data output is 1.9 Mbytes/sec. The fastest that the data can be transferred from the frame-grabber to the system board is limited by the DMA rate of 0.380 Mbytes/sec. This is a factor of five too slow and does not include time for processing.

The video camera system for collecting data images in the experimental setup is not optimal for a real time system. Effort would be required to match the position of images in wood from frame to frame of the camera. A more practical setup for a real time system would be to employ a line-scan CCD array camera with the the scan rate triggered by the wood feeder to have a calibrated distance coordinate. Such a system is described in more detail in chapter 12.

5.4 CONCLUSIONS

The equipment used is sufficient to demonstrate the methods of automated inspection in an off-line situation. The most important limitation of the equipment is the lack of absolute gain control in the camera. This directs the research of identification algorithms toward those that are not dependent upon absolute light levels. This is considered to be a positive influence as it leads to the development of more robust algorithms.

6. FEATURE DETECTION USING TONAL MEASURES

The automated visual inspection of timber is constrained by a number of inter-related factors. The most dominant of these is the speed with which the inspection must be performed. Computers are limited in the speed with which they can operate and so considerable effort must be expended to reduce the computational burden. The requirement is for fast algorithms to perform the neccessary computations and an analysis structure that restricts the amount of computation that must be performed. Feature detection is the method by which the amount of computaion is reduced by restricting analysis to the areas of the board that contain features.

It is proposed to segment the image into small local areas, each of which will be examined and classified as being either a feature local area or a clear local area. This classification is performed using a number of measures of the local area to catagorise it into the two classes of the dichotomy. The clear areas can then be dispensed with as they no longer contribute to the grading of the wood and the processing is then confined to the feature local areas. The feature local areas are combined, if they are adjacent and form part of the same feature, into larger feature areas. These feature areas are then assumed to cover the entire feature and the feature is then extracted and identified from within the feature area boundary. This later processing is dependent upon the ability of the feature detection level of processing to correctly identify the feature local areas to a suitable resolution.

6.1 DESCRIPTION OF TONAL MEASURES

In order to place a local area into a catagory or class a decision must be performed based on aspects of the area that

distinguish it from areas in another class. The distinguishing aspects that are calculated for an area are called feature measures and the process of classifying a set of local areas into classes is pattern recognition.

The success of a pattern recognition system is crucially dependent upon the selection of the feature measures. They must represent distinguishing aspects of the features that are to be recognised. Clear wood is characterised by growth rings which appear as longitudinal dark bands. The growth rings are present in all the wood and are, at times, more prominent than the features: this is a dominant aspect of the problem. These growth rings vary in appearance from wide light bands to narrow dark lines depending on the angle at which they are cut and the quantity of resin within them which is a natural variable. They are commonly distorted around intergrown knots and spike knots but are not so disturbed by the presence of encased knots. What is needed is a computationally cheap method to separate the features from the growth rings in order to confine further processing to the identification of the features. To this end a number of calculable measures of the features must be chosen that uniquely and succinctly describe the features of interest. The problem then is to choose these feature measures.

The features themselves vary markedly. Pin knots can be less than 10mm in diameter and can vary in appearance from a sharp, dark dot if it is encased to a barely discernable variation in the growth rings if it is sufficiently intergrown. Large knots are over 35mm in diameter and are usually a dark red colour but can sometimes have a lighter interior. Pith is a dark strip about 10mm wide which tapers at each end and may be over 500mm long. Subjectively, a feature in wood is prominent because it is different from the background due to intensity, shape, texture, edges, and size. It is usually darker, and the edges sharper, than the background of growth rings. This subjective appraisal acts as a guide to the choice of feature measures.

Tatari (1987) presents a table of correspondence between defect classes and inspection methods that strongly suggests the

use of parameters of the grey level histogram i.e. the tonal measures of the mean, variance, skewness, and kurtosis. Conners et al (1983) find that using tonal measures to classify local areas into a dichotomy of clear and feature (defect) areas is accurate to 96.4%. These results suggest that tonal measures would be appropriate to quantify some of these subjective descriptions.

The grey level histogram is described in chapter 5. To repeat, the digitised image is made up of an array of pixels (picture elements), each of which can have a grey scale value from 0 to 255 (black to white). For a given area of an image a histogram is constructed by summing the pixels of each grey level to form an array that represents the distribution of grey levels in the area. Statistics of this array convey fundamental information of the structure of the area. The tonal measures that are used are the mean, variance, skewness, and kurtosis of the grey level histogram. These are:

1) mean

$$\mu = \sum_{l=0}^{L-1} lP(l);$$

2) variance

$$\sigma^2 = \sum_{l=0}^{L-1} (l - \mu)^2 P(l);$$

3) skewness

$$s = \sum_{l=0}^{L-1} (l - \mu)^3 P(l)/\sigma^{3/2};$$

4) kurtosis (flatness)

$$k = \sum_{l=0}^{L-1} (1 - \mu)^4 \, P(l)/\sigma^2;$$

where P(l) is the estimated probability of grey level l occurring within a local area and L is the total number of possible grey levels in the image (256).

P(l) is calculated by dividing the number of pixels of the grey level, l, by the total number of pixels in the local area. This gives the grey level probability distribution as distinct from the grey level histogram which is an array of integer values. The mean is simply the average grey level of the area and is a measure of the overall brightness of the area. The standard deviation is the square root of the variance and describes how spread out the brightness values are from the mean. A small standard deviation implies that there are not many pixels that are brighter or dimmer than the mean value. A large standard deviation implies that the pixels cover a wide range of brightness values, with very dark areas and/or very bright areas. The easily visualisable aspect of standard deviation is that the range of grey levels from $(\mu - 3\sigma)$ to $(\mu + 3\sigma)$ will include 99.7% of the pixels in the image if they are distributed in a normal distribution.

Skewness is a measure of how much the histogram varies from a symmetrical shape. An image of uniform brightness will have zero skewness as will an area that contains the same number of pixels for each grey level above the mean as below it. An otherwise uniform area that contains a small number of darker pixels will have a negative skewness while a number of brighter pixels will give it a positive skewness. It is visualisable as the direction of the "tail" of the histogram. A dark tail is a negative skewness and a light tail is a positive skewness.

The kurtosis is a measure of the flatness of an image. It can be visualised as the deviation of the histogram from a

normal distribution. A smooth uniform area will have a low kurtosis, that is, it will be relatively flat. A small dark or light blob in the image will be represented by a small peak in the histogram that is some distance from the central mean peak. It will increase the value of the kurtosis and the image will be seen to be relatively unflat. A similar small peak or tail in the histogram can be due to a number of individual pixels with values far from the mean scattered throughout the image. The visualisation of the image as unflat is reinforced.

The usefulness of these measures to identify feature areas from the background of annual growth rings can be predicted in a qualitative manner. The mean will be important because the dominant aspect of many features is that they are darker than the background wood. This is moderated by two factors. The first is that the background mean level is variable. This is apparent in the difference between heartwood and sapwood noted in chapter 2. The second is the effect of the gain control of the camera which adjusts the gain of the camera according to the overall brightness level in the field of view. This is an uncontrollable factor with the camera used. This may limit the usefulness of the mean as a measure to distinguish feature areas.

A feature is recognisable as a change in the brightness of an area. These changes, the small dark pin knot, the section of pith, the edge of a large knot, will increase the variance and kurtosis of a local area and so make these measures useful discriminators. Many feature areas will contain a small portion of a dark feature on a lighter background and so will have a negative skewness but some will contain a small portion of lighter background and be comprised almost wholly of the darker feature. These latter local areas will have a positive skewness. Local areas that contain all background will have near zero skewness, a relatively high mean and low values for variance and kurtosis.

To quantify these statements the tonal measures are calculated for images in the FEAT library but first a decision is made on the size of the local area that is used.

6.2 DIVIDING THE IMAGE INTO LOCAL AREAS

The proposed strategy is to divide each image of the plank into smaller local areas and to class these local areas into feature or clear areas based upon measurements of the area information. Since the total area of the board must be inspected for the presence of a feature, it becomes necessary to segment the image area into suitably sized areas for processing. The determination of a suitable size for the local areas is required before proceeding to further examine of the issue of feature measures. The optimum size for these local areas is assessed in this section.

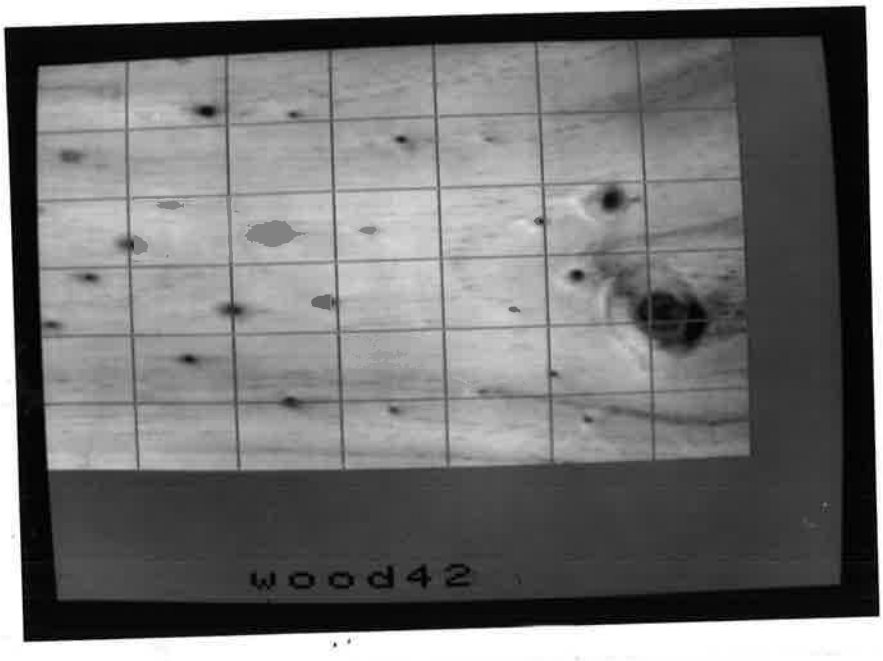

Figure 6.1 The division of WOOD42 into 64 * 64 local areas.

An example of the division of an image into local areas is presented in figure 6.1. The image is divided into non-overlapping areas of size 64 * 64 pixels. The small knot on the right of the image falls into four neighbouring local areas while some local areas contain two features (in this case needle trace).

An intuitive analysis suggests a number of general points. A large local area will tend to obscure the smaller

features within the whole. Small pin knots, splits and bark will not be able to exert a significant influence to contribute toward the classification of the local area as a feature area. There is also the problem that a large local area may contain a number of features within its extent which may lead to difficulties of separation. This is a problem in some methods of feature extraction but not others and will be highlighted at the appropriate time.

Conversely, a small local area will not be able to distinguish the larger characteristics of some features. Areas such as the dark interior of a large knot will appear as uniform while the busy areas of the edges of features will all be defined. This may present problems associated with connecting the feature local areas together into larger feature areas. There is also the computational cost associated with small local areas. To calculate the grey level histogram of a local area requires the reading of each pixel once, and a single addition to increment the array element for the pixels grey level. The histogram array is 256 elements long corresponding to the 256 grey levels of the image. Further calculation to determine the mean, variance, skewness and kurtosis only requires reference to the histogram. So, while each pixel must be read once for the entire image regardless of the local area size, the histogram array must be read once for each of the tonal measures for each local area. The smaller the local area the more histogram arrays must be read for the same area size and so the greater the computational burden.

Before going further an alternative to the choice of rectangular local areas needs to be considered. Since, in a production system, the data would be from a linear array camera, would it not be expeditious to use a local area in the form of a strip the width of the board, this being the total output from a number of line scans? This has the problem that the wood will have a number of widths, constant during a run to be sure, but varying from 290mm to 42mm on different runs. This could be overcome by changing the number of line scans that make up a local area to ensure that each local area is of identical size

but would lead to inefficiencies in the definition of feature areas. For example, any feature would be enclosed in a feature area (the group of connected local areas) that is the width of the board and would comprise of a large proportion of featureless area. For this reason it is decided that it is more efficient to use rectangular local areas. (The areas are actually a square array of pixels but the aspect ratio of the camera means that the square array of pixels represents a rectangle of side ratio 6:5 with the longer side in the direction of the length of the board.)

Local areas of three sizes are examined using the FEAT library of images. The local areas are square areas of 16, 32, and 64 pixels per side - corresponding to areas with side length of approximately 12.5mm, 25mm, and 50mm on the wood board. Each of the four tonal measures is examined separately for its usefulness in identifying a feature/non-feature threshold. The criterion used to determine the effectiveness of each measure is its ability, uniquely and sparingly, to define a feature. For a single measure, a particular threshold value will define a number of local areas as feature areas. For a particular feature there exists a range of thresholds for which the local area(s) that encompass the feature are classified as feature local areas yet neighbouring non-feature areas are not defined as feature local areas. This is the case of the feature area being uniquely and sparingly defined. A threshold value above this range will fail to classify one or more feature local areas correctly and so the feature will not be fully defined. A threshold value below this range will incorrectly classify clear local areas as feature local areas and the size of the feature area will become extended past the boundary of the feature. This leads to the inclusion of a lot of clear area in the feature area. Stating with a very large threshold value, the upper limit of the threshold range is identified when the local areas encompass the extent of the feature and the local areas covering the same feature display connectivity so that the whole feature is described. As the threshold is lowered, more areas become identified as feature

areas and the lower limit of the threshold range is defined as that which does not include local areas extending the size of the group unnecessarily past the boundaries of the feature or connecting the group to a neighbouring feature group.

The position of the feature in a local area is noted to have an effect on the feature measure thresholds, but this effect is considered to be randomly distributed for the large sample of the FEAT library. For example, some features are fully contained within a single local area and so have a high upper threshold, whilst others are on a corner, overlapping into four local areas, and need a low upper threshold to provide connectivity that defines the whole feature.

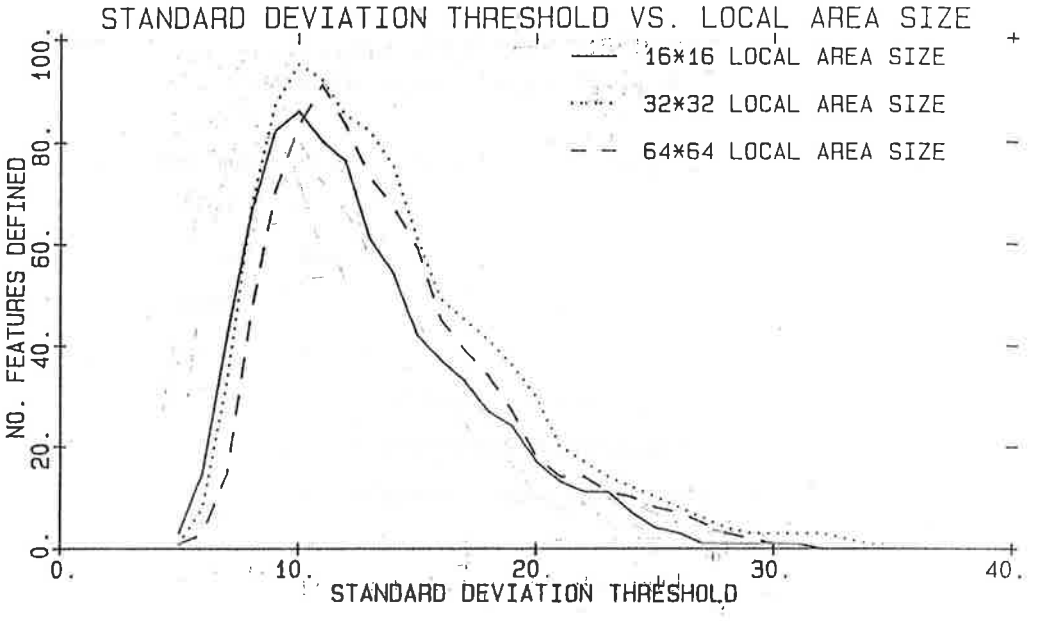

Figure 6.2 The ability of standard deviation to define 118 features with three local area sizes.

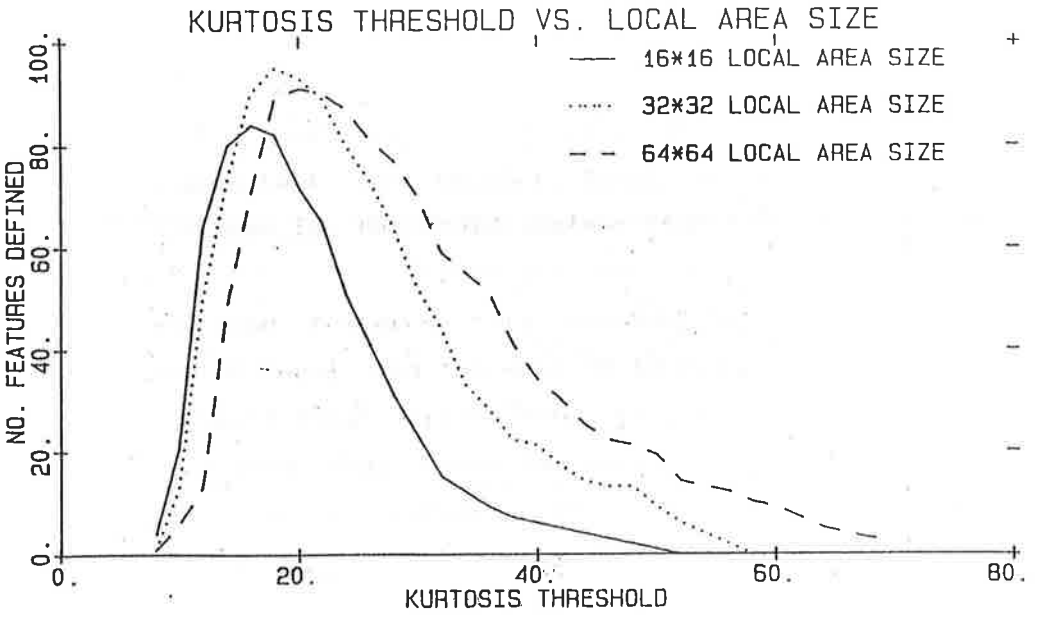

Figure 6.3 The ability of kurtosis to define 118 features with three local area sizes.

The most striking measure of the features in wood is the mean threshold, as pith and most knots are darker than the surrounding wood. The mean, however, is susceptible to changing lighting conditions and the changeable nature of the wood as mentioned previously. The higher order tonal measures are useful because they are isolated from these variations. The results for standard deviation and kurtosis thresholds are shown in figures 6.2 and 6.3. No meaningful threshold can be deduced from the skewness data alone.

For the selection of a suitable standard deviation threshold, the local area size is not critically important. A peak extraction of 95/118 (80.5%) of the features is achieved for a standard deviation threshold of 10 with a 32*32 local area size, and a similar extraction of 91/118 (77.1%) is obtained for the larger local area size of 64 * 64. The kurtosis threshold exhibits a maximum extraction of 100/118 (84.7%) with the 32 * 32 local area size with a drop in performance to 93/118 (78.8%) extraction with the larger 64 * 64 local area size. For both measures, the drop in performance is more marked with the smaller

16 * 16 local area size droping to 86/118 (72.9%) for both standard deviation and kurtosis. The kurtosis thresholds exhibit different rates of change at thresholds greater than the maximum extraction for the different local area sizes. This can be interpreted to mean that the smaller local area sizes are more sensitive to a change in threshold values than the larger local areas.

Because wood features vary considerably in size and type, the choice of local area size needs to balance a number of considerations. Small local areas find the smaller features in the wood and, provided that there is connectivity between them, also identify the larger features. They do, however, carry a large computational penalty in the calculation of the feature measures. Large local areas speed the detection process, but tend to overlook the small fainter features, such as intergrown pin knots; however, these faint features often have no influence on, or can be tolerated by, the grading process and so this type of error is not crucial. A significant effect of large feature areas is that they can cause individual nearby features to be grouped within a single feature area. Failure to detect features is an important error and is more likely to happen with a small local area size.

6.3 CONCLUSIONS ON THE OPTIMUM LOCAL AREA SIZE

It is concluded that the local area size needs to be of the order of the feature size. The results of the previous section indicate that a local area of size from 25mm to 50mm gives the best results. Features vary considerably in size but only a small proportion are larger than 50 mm. The 64 * 64 local area size corresponds to a size of approximately 50 mm on the wood surface and is used in chapter 8.

The use of tonal measures to detect features shows that these measures are useful but that no single measure provides an adequate description alone. The next chapter details the use of pattern classification as a means to combine a number of measures to detect local areas that contain portions of features.

7. PATTERN CLASSIFICATION

7.1 THE NEED FOR PATTERN CLASSIFICATION

The results obtained when choosing a suitable local area size indicate that using the threshold of a single feature measure is inadequate to perform the dichotomy into feature and non-feature areas. They do, however, indicate that the tonal measures of standard deviation and kurtosis are suitable measures for performing the dichotomy discrimination. The mean and skewness may prove useful if they can be combined in a suitable way. What is really needed is a means of combining a number of feature measures so that they can all contribute to the determination of the classification of a local area.

It is recognised at this point that the problem is one of pattern classification. The classification can be into any number of classes. In this study the two classes are feature local areas and clear (non-feature) local areas. Each local area is a sample with a number of measures - until now the four tonal measures - that enable it to be recognised as in one of the two classes. Pattern classification has many uses and a variety of different methods have been developed for different conditions. It is a statistical method for determining patterns in multi-dimensional data. The conditions and method appropriate for this study, based on the type of data and the number of classes, is described in the following text.

7.2 PATTERN CLASSIFICATION METHODS

The aim of pattern classification is to make a decision in a condition of uncertainty. The decision in this study is whether a local area is clear or contains part of a feature within it. The uncertainty is present because there is no single measure of the local area that can be used unconditionally to

identify the area as belonging to one class or the other. To make a decision a rule must be deduced based on certain measurements, the feature measures. The stages in the derivation of a decision rule are indicated in figure (7.1). The physical system that is to be analysed is a piece of wood.

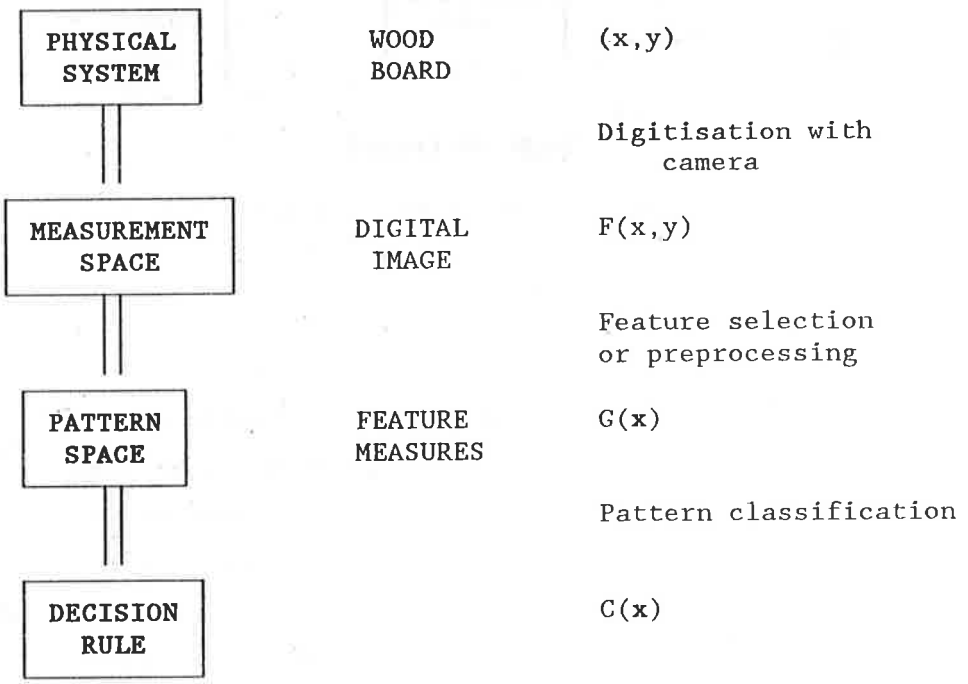

Figure 7.1 The stages in the derivation of a pattern classification decision rule.

The ultimate goal of this entire study is the classification of wood by visual appearance. For this reason, the measurement of the physical system is made by a camera sensitive to the optical band of the electromagnetic spectrum. The camera is panchromatic across this band to give a similar spectral response to the human eye but represents only the intensity of the light and not the colour. This image data is stored as an array of pixels in eight bit planes, giving 256 grey intensity levels and the array is of size 512 * 512. This forms the measurement space, F(x,y).

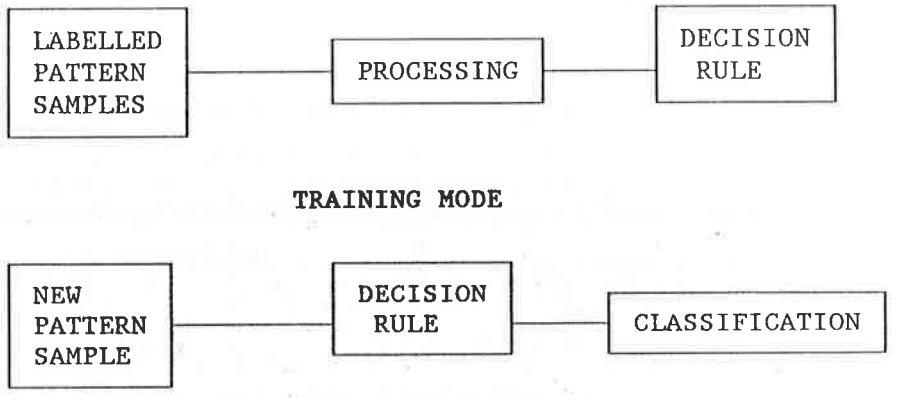

Figure 7.2 Training and operating modes of a classification system.

Already the transformation to measurement space has reduced the information content of the image. Depth information is no longer available and neither is colour information. The transformation to pattern space requires the selection of suitable features that contain sufficient information to classify the image. Pattern space must also be of relatively low dimension to enable the calculation of a meaningful decision rule. Once the problem is characterised by the transformation to a meaningful pattern space, the abstraction to a decision rule is performed. This process is shown in figure (7.2). A number of labeled samples are processed to form a decision rule in a training, or learning mode. This rule is then used to classify unknown samples in a run or operating mode.

The two main problems in pattern classification are characterisation of the measurement space and abstraction into a decision rule. These aspects are the subject of the following sections.

7.3 THE CHARACTERISATION PROBLEM - FEATURE SELECTION

In order to make a sensible transformation from measurement space to pattern space features that encapsulate salient aspects of the image must be chosen. These must be of <u>low dimension</u> to make the problem tractable. They must contain

sufficient information to describe the problem at hand. They must display a geometrical consistency, with differences in the quality to be recognised being reflected as a distance in pattern space. The features must be consistent throughout the samples. Let us examine these criterion individually with reference to the problem of wood inspection.

7.3.1 DIMENSIONALITY

The features defining the pattern space must be of minimal number. The number of samples needed to ensure a representative distribution of class samples increases exponentially with dimension. For example, if it is believed that two samples are enough to describe a class in one dimension then to describe two dimensions four samples are needed. Three dimensions would require eight samples and n dimensions would require 2^n samples to provide a very sparse distribution. Eleven dimensions would require $2^{11} = 2048$ samples per class to provide enough information to provide a description.

The four features used in the two class problem of detecting feature areas in wood are the tonal measures of mean, variance, skewness, and kurtosis. These require a minimum of 16 samples of each class when used together to provide a decision rule based on the most meagre data.

It can be visualised intuitively that if the sample size is too small and/or dimensions containing random data are introduced the division into classes may well be an artifact. Foley (1971) provides theoretical and experimental evidence that a value of M/n of at least 3 to 5 is necessary to prevent random data appearing significant, where n is the number of classes and M is the number of samples per class. For the FEAT training set of 686 local areas there are 159 feature areas and 527 clear areas. With 11 dimensions M/n = 159/11 = 14.5. If some of the measures are not relevant then the intrinsic dimensionality is less than 11 and so the value of M/n will be greater. Consequently, the sample size can be considered adequate.

While a large M/n enhances the validity of the derived decision rule, a small M/n does not necessarily mean that the rule is meaningless. If only 3 of the 11 dimensions used in the derivation of a rule contain useful information and the other 8 are essentially random data then the n used does not reflect the true situation. The value of n used in M/n should reflect the intrinsic dimensionality. The fact that the problem is posed in terms of n variables is not sufficient to determine sample requirements. However, the error induced by having redundant measures will increase the value of M/n and so improve the validity of the results.

If a measure is irrelevant its coefficient will tend to zero as the iteration proceeds. If two measures are correlated and contain information there is no way to tell, from the value of the coefficient, that the correlation exists. The only way to detect this correlation is to derive a decision boundary using only one of the two measures and then repeat the derivation with the other measure. Correlation is revealed by the similarity in the misclassified samples.

7.3.2 SUFFICIENT INFORMATION

If the measures used to describe a problem do not contain the essential information for the classification then a satisfactory decision rule will be difficult, if not impossible, to derive. The information to be extracted from the measurement space is the relevant information to discriminate the classes.

In the development of feature measures to describe the essential aspects of an image, there is a level of interaction required by the researcher to clarify these measures. A simple, if tedious, method is to start with a minimum of features that guarantee a reasonable decision rule, such as the four tonal measures in the case of wood images. These are tested for their effectiveness by trying to derive decision rules using various combinations of subsets of the four measures. A minimum number of misclassifications is achieved using all four indicating that they all contribute significant information to the decision rule.

Using these four tonal measures as a basis upon which to build, further measures are investigated individually as a fifth measure to determine its effectiveness. Only if the fifth measure improves the decision rule is it regarded as effective. This requires the calculation of a decision rule for each new measure that is to be examined and then the testing of this rule against a test set of samples.

7.3.3 CONSISTENCY OF FEATURES

Geometrical consistency refers to the difference between classes in measurement space being reflected as a difference in pattern space. For example, dark features will give a feature class local area a low value for the mean, while clear local areas will have a higher mean value. Error will arise if an intergrown knot has a mean value similar to that of a clear area. To detect the difference between intergrown knots and the growth rings a feature measure is needed that shows a large value difference between the two cases.

The features must also be consistent throughout the sample set. If a different species of wood is to be inspected that has a darker overall appearance then the mean coefficient calculated for the former will produce erroneous classifications in the latter.

7.4 THE ABSTRACTION PROBLEM

Once a pattern space has been selected that is considered to represent the essential information of the images the next problem is the determination of a decision rule that can be used to classify new images such that:

$$C(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} \in S_0 \\ 1 & \text{if } \mathbf{x} \in S_1 \end{cases} \qquad (7.1)$$

where C is the decision rule; S_0 is the subset of clear areas; S_1 is the subset of feature areas; S is the set of all local areas;

and $\mathbf{x} = (x_1, x_2, x_3, \ldots, x_n)$ is the vector that defines the sample in the pattern space of n dimensions.

The decision function is calculated from a set of labeled samples called the **training set** using an iterative technique. The decision function generated is then tested on a second set of labeled samples called the **test set**.

The selection of the training and test sets is arbitrary in the sense that no particular distinction is made between them. Both sets are chosen to contain approximately the same number of each type of feature, ie, large, medium, small, and pin knots, plus pith. Due to the varying image size and the random nature of the local area subdivision, the two sets contain different numbers of local areas. This has no effect on the validity of the results as long as the training set is sufficiently large.

7.4.1 NORMALISATION AND DISTANCE MEASURES

A crucial concept in pattern recognition is that of distance. It is assumed that the closer two points are in pattern space the more similar are the patterns represented by the two points. In order for distance measures to have any meaning the feature measures need to be normalised to a similar range and spread.

Normalisation of a feature measure can be done in two ways. The first is to normalise the measures based upon the range of each measure. Let:

$$a_k = \max\{y_{1k}^{(i)}, y_{2k}^{(i)}, \ldots, y_{M_i k}^{(i)};\ i = 1, 2, \ldots, N\}$$

$$- \min\{y_{1k}^{(i)}, y_{2k}^{(i)}, \ldots, y_{M_i k}^{(i)};\ i = 1, 2, \ldots, N\} \qquad (7.2)$$

where $k = 1, 2, \ldots, n$. a_k is the range of the kth feature measure over all the sample points and the normalised variables become

$$(x_k - a_k/2)/a_k. \qquad (7.3)$$

The second method is performed using the mean and variance of the samples in the training set.

$$\bar{x}_k = \frac{1}{M} \sum_{i=1}^{N} \sum_{j=1}^{M_i} x_{jk}^{(i)} \qquad (7.4)$$

where $M = M_1 + \ldots + M_N$ is the sum of all the samples in each class ($N = 2$ in this study) and $k = 1, 2, \ldots, n$, and n is the number of feature measures. This gives the mean of each feature measure for the whole set of samples. The variance of each feature measure is then:

$$\bar{\sigma}_k^2 = \frac{1}{M} \sum_{i=1}^{N} \sum_{j=1}^{M_i} (x_{jk}^{(i)} - \bar{x}_k) \qquad (7.5)$$

and the normalisation measure is:

$$\frac{(x_{jk}^{(i)} - \bar{x}_k)}{\sigma^2} \qquad (7.6)$$

It has been shown that the latter method is optimum (Sebestyen (1962)). The distance between two samples can now be described by the Euclidean distance which has the advantage of having a derivative at all points.

7.4.2 DECISION BOUNDARY

In this two class problem the linear decision boundary between the two classes is a (n-1)-dimensional hyperplane in n-dimensional pattern space. This can be visualised as a line in 2-space or a plane in 3-space. A decision boundary hyperplane is defined by the coefficients of the weight vector:

$$\mathbf{w} = (w_1, w_2, \ldots, w_n, w_{n+1})$$

By interpreting $\mathbf{x}$ as an augmented vector such that

$$\mathbf{x} = (x_1, x_2, \ldots, x_n, 1)$$

then the hyperplane is the locus of points satisfying the equation

$$w_1x_1 + w_2x_2 + \ldots + w_nx_n + w_{n+1} = 0 \qquad (7.7a)$$

or

$$\mathbf{w} \bullet \mathbf{x} = 0 \qquad (7.7b)$$

For convenience it shall also be defined that:

$$\omega = (w_1, w_2, \ldots, w_n) \qquad (7.7c)$$

If a sample $\mathbf{x}$ is on one side of the boundary it is in class S_0; if on the other side, in class S_1. If it is on the boundary it has to be classified arbitrarily. So, we get the decision rule in terms of a linear discriminant function:

$$C(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} \bullet \mathbf{w} \geq 0 \\ 1 & \text{if } \mathbf{x} \bullet \mathbf{w} < 0 \end{cases} \qquad (7.8)$$

More complex decision boundaries such as the piece-wise linear decision boundary may be calculated using alternative techniques (Nilsson (1965), Meisel (1972), Patrick (1972), Bow (1984)) but these are not included in the present study. The next stage of the abstraction problem, once the form of the decision boundary has been decided upon, is the method for the calculation of the optimum decision hyperplane.

7.4.3 DATA TYPES

The method that is used to determine the optimum linear decision boundary is determined by the type of data that is

available. To review the different input data types mentioned in the literature survey, recall that there are four types of data knowledge:

(a) The functional form of the conditional density function, $p(x|S_i,\theta)$, is known for all the classes to within the specification of a set of parameters θ. For example, it may be known that the pattern vectors of each class are distributed in a Gaussian distribution with unknown mean and covariances.

(b) The parameters of the conditional density functions are also known.

(c) Sample patterns of known classification are available. This is the situation of a training set of labelled samples.

(d) The samples have no classification. This is the situation of learning without a teacher.

The status of the present study is data of type (c), that of labelled samples. A training set of images is divided into local areas, the feature measures of each local area forming a sample in pattern space. Each local area is given a class as either a feature or a clear area. Data type (d) does not apply. If data type (a) and (c) apply then methods such as Bayes criterion can be applied as is done by Conners (1983). The class conditional probabilities are considered not to be known or assumed in the present study and this leads to the adoption of deterministic methods of finding a linear decision boundary.

7.5 CHOOSING A CLASSIFICATION METHOD

The pattern classification method used to analyse the present work was chosen for a number of reasons. Firstly, the class of each local area is specified beforehand to make up what is called the training set. The presence of a labeled set of samples means that deterministic methods are used. This takes the form of an iterative procedure to find the minimum of a cost function.

Secondly, the distribution of the samples in pattern space is not known in terms of a functional form, or the parameters thereof. A non-parametric method escapes having to

make assumptions of the functional form of the conditional densities of each feature measure but it requires that a suitable choice of cost function be decided upon.

Thirdly, a linear technique is chosen for simplicity over piece-wise linear or more complex forms. The linear decision boundary that is generated that separates the two classes in feature measure space is then a line in two dimensions, a plane in three dimensions, and a linear hyperplane in higher dimensions. This makes the assumption that each feature measure varies in a linear manner from one class to the other.

The method used is described fully by Meisel (1972).

7.5.1 THE PARAMETERS OF THE COST FUNCTION

An expression for the average loss over the labelled samples is given by the approximate risk:

$$\tilde{R} = \frac{1}{M} \sum_{i=1}^{N} \sum_{j=1}^{M_i} L[C(\mathbf{x}_j^{(i)}), i] \qquad (7.9)$$

This is the cost function to be minimised where R is the overall measure of the separability of the classes. It is in terms of an as yet undefined loss function, $L[C(\mathbf{x}),i]$, the loss associated with the classification of a sample $\mathbf{x}$. The loss function involves $C(\mathbf{x})$ which in turn is specified by the parameters $\mathbf{w}$.

The cost function to be minimised must meaningfully express the "goodness" of a solution while also being solvable. The minimum of the cost function is to be found using a gradient technique so it is desirable for it to have a continuous derivative. It is preferable to have a cost function with only a single minimum or equal minima to avoid getting stuck at local minima. It is also desirable if the cost function is convex.

The most obvious cost function is to use the number of samples misclassified. This has many problems, the first of which

is the non-uniqueness of the solutions. There are an infinite number of solutions for a particular number of misclassifications. This leads on to the fact that the cost function is not smooth but moves in discrete jumps from one quantity of misclassifications to another. At points of **w** other than at a jump the derivative of the cost function with respect to **w** is zero. This makes it impossible to determine a suitable direction for the movement of the hyperplane.

A more satisfactory form of the cost function is one in which the loss function is based upon the distance of the samples from the hyperplane. It is here that the importance of an adequate normalisation procedure becomes evident. Only by normalisation can the distance between samples in pattern space be expected to have meaning.

At this point it is necessary to conform to the conventional notation. The set **S** of all samples is made up of the set of clear samples, S_0, and the negative of the set of feature samples, S_1, where each sample is an augmented vector. The picture in augmented pattern space is that the hyperplane being sought passes through the origin with all the samples in **S** on the positive side. This notation is computationally convenient but it is easier to visualise the unaugmented pattern space. For this reason the mathematical explanation is in augmented terms but the figures and descriptions are in unaugmented terms.

For the purposes of optimisation the hyperplane has a dead zone, dz, on either side. The dead zone defines which samples contribute toward the minimization of the cost function by describing a hyperplane $\mathbf{w} \bullet \mathbf{y} = dz$ that is parallel to the hyperplane of the decision boundary $\mathbf{w} \bullet \mathbf{y} = 0$ and a distance $dz/||\omega||$ away. The larger the value of dz the more of the samples will contribute toward the cost function. A sample can be on the correct side of the hyperplane yet still be on the wrong side of the dead zone and so still contribute. The loss function applies from the dead zone so a change in the dead zone will change the loss associated with a particular sample. The distance of a sample **x** from the dead zone hyperplane is:

$$|\mathbf{w} \bullet \mathbf{x} - dz| / ||\omega|| \qquad (7.10)$$

This distance measure is the Euclidean distance and changes with scale changes of the sample, **x**. This is a reason for a satisfactory normalisation procedure to ensure that a measure with a large range does not reduce the influence of the other measures. In the case of separable classes a suitable value of dz will place the decision boundary in the space between the two classes in an optimum position among the non-unique solutions.

The effect of dz can be reduced by multiplying **w** by a constant. This means that if, as the iteration proceeds, the coefficients of the weight vector **w** increase then the relative effect of dz is reduced. A reduced dead zone means that fewer samples will contribute to the updating of the hyperplane.

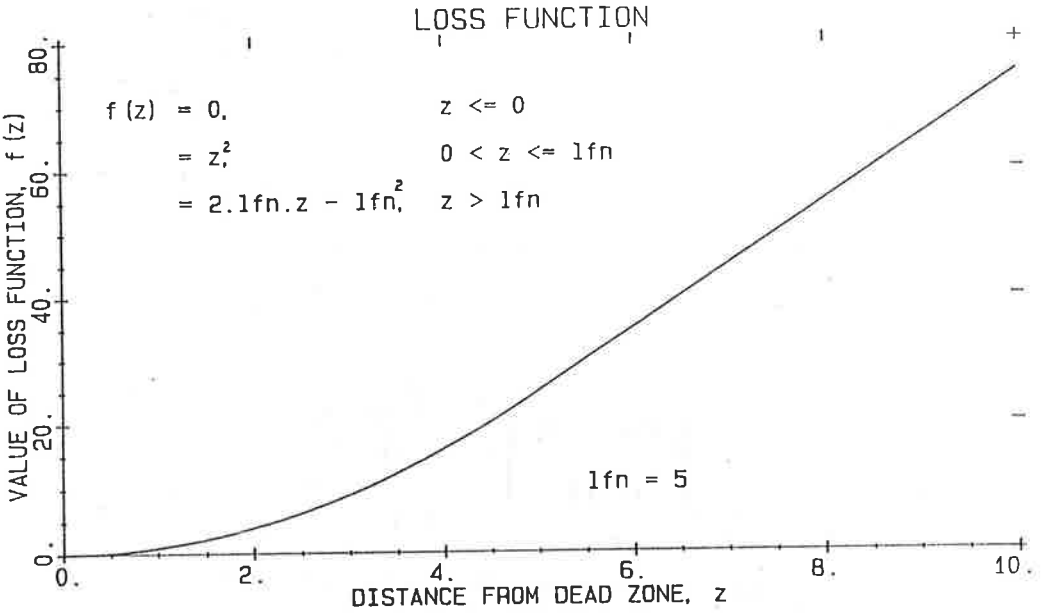

Figure 7.3 Form of the loss function.

The loss function describes the loss associated with each sample that contributes toward the iteration of the cost function. It is chosen to have continuous derivatives to aid the iteration calculations. The form chosen is a square function with a linear extension. It is based on the distance of the sample from the hyperplane of the decision boundary. This distance is squared to provide a measure of the loss. The loss function

parameter that can be varied is lfn which determines the distance from the dead zone at which the loss function makes the transition from a square function to a linear function. The greater the value of lfn the steeper will be the linear portion of the loss function and so the greater the loss attributed to a sample that lies more than lfn from the dead zone. The purpose of the loss function is to give a heavier weight to the samples that are a long way from the dead zone (very misclassified) and a lesser weight to those nearer the dead zone hyperplane. The loss function used is illustrated in figure 7.3 and has the form:

$$f(z) = \begin{cases} 0, & z \le 0 \quad \text{zero loss, outside dead zone.} \\ z^2, & 0 < z \le \text{lfn} \quad \text{square function,} \\ 2.\text{lfn}.z - \text{lfn}^2, & z > \text{lfn} \quad \text{linear function.} \end{cases} \tag{7.11}$$

Having defined the loss function, L[C(**x**),i], as a function of the distance of the sample from the dead zone the cost function can now be expressed as:

$$\tilde{R}(\mathbf{w}) = \frac{1}{M} \sum_{j=1}^{M} f\left[\frac{dz - \mathbf{w} \bullet \mathbf{x}_j}{||\omega||} \right] \tag{7.12}$$

This uses the actual distance of a sample from the dead zone but since (dz - **w•x**) is proportional to the distance the gradient function can be simplified by using this instead. The cost function becomes:

$$\tilde{R}'(\mathbf{w}) = \frac{1}{M} \sum_{j=1}^{M} f(dz - \mathbf{w} \bullet \mathbf{x}_j) \tag{7.13}$$

A gradient technique is used to update **w** and minimise $\tilde{R}'(\mathbf{w})$:

$$w^{(i+1)} = w^{(i)} + \epsilon_i \widetilde{\nabla R}'(w^{(i)}) \qquad i = 0,1,2,3, \ldots \qquad (7.14)$$

where ϵ_i = gradient increment.

A convenient formulation of the gradient function can then be expressed by defining:

$$v(i) = \frac{1}{M} \sum_{j \in J(w^{(i)})} f'(dz - w^{(i)} \bullet x_j) x_j \qquad (7.15)$$

where $J(w^{(i)}) = \{j | w \bullet x < dz, \; x \epsilon S\}$

and

$$f'(dz - w^{(i)} \bullet x_j) = \begin{cases} 2(dz - w^{(i)} \bullet x_j) & \text{if } (dz - w^{(i)} \bullet x_j) \leq \text{lfn} \\ 2.\text{lfn} & \text{if } (dz - w^{(i)} \bullet x_j) > \text{lfn} \end{cases} \qquad (7.16)$$

Using this notation, equation (7.14) becomes:

$$w^{(i+1)} = w^{(i)} + \epsilon_i \, v(i) \qquad (7.17)$$

v(i) is proportional to a weighted average of the samples defined by the dead zone, with the weight of a sample a function of the distance from that point to the hyperplane. The sample points furthest from the hyperplane have the greatest influence as defined by the loss function.

Now, ϵ_i can be calculated to be the increment by which the weight vector is moved toward the minimum of the cost function 'down' the gradient:

$$\epsilon_i = \lambda \frac{|w^{(i)} \bullet v(i) - dz|}{||v(i)||^2} \qquad (7.18)$$

If the step is unity (λ=1) the hyperplane is moved exactly to the separating postion inferred by the gradient. In practice, the two classes overlap and no separating position

exists and it is found that a small step, $\lambda \ll 1$, is necessary to prevent oscillations in the iteration.

The loss ratio is a factor which applies to the loss function. It is a way of attributing a different loss to misclassifying a class 0 sample as a class 1 as opposed to the converse. This is useful because the loss is greater if a feature area is wrongly classified as a clear area than if a clear area is classed as a feature area. The reason for this is that clear areas are not kept for further processing and so if the area contains a feature that feature will not be detected. If a clear area is classed as a feature area it can be discarded after further processing has determined that it is clear. The value of the loss ratio, lossrat, is applied in the following manner:

$$f'(dz - w^{(i)} \bullet x_j) = \begin{cases} f'(dz - w^{(i)} \bullet x_j) & \text{if } x_j \in S_0 \\ \text{lossrat} * f'(dz - w^{(i)} \bullet x_j) & \text{if } x_j \in S_1 \end{cases} \tag{7.20}$$

In equation (7.14), all the samples are used in each iteration to obtain the values for the new hyperplane: this is called the many-at-a-time approach. The one-at-a-time algorithm is an alternative where a single point is used for each iteration and the samples are arranged in a sequence $x^1, x^2, x^3, \ldots$

This one-at-a-time algorithm is:

$$w^{(i+1)} = \begin{cases} w^{(i)} + \epsilon_i f'[dz - w^{(i)} \bullet x(i)]x(i) & \text{if } w^{(i)} \bullet x(i) < dz \\ w^{(i)} & \text{if } w^{(i)} \bullet x(i) \geq dz \end{cases} \tag{7.19}$$

where $x(i) = x_{i \bmod M}$ and the samples are cycled through continuously.

7.6 VISUALISATION OF PATTERN SPACE

A greater appreciation of the method of pattern classification can be gleaned if one can visualise the pattern space in which the algorithm works. The following description of pattern classification starts in the simple case of two dimensional pattern space and then extends the description into higher dimensions.

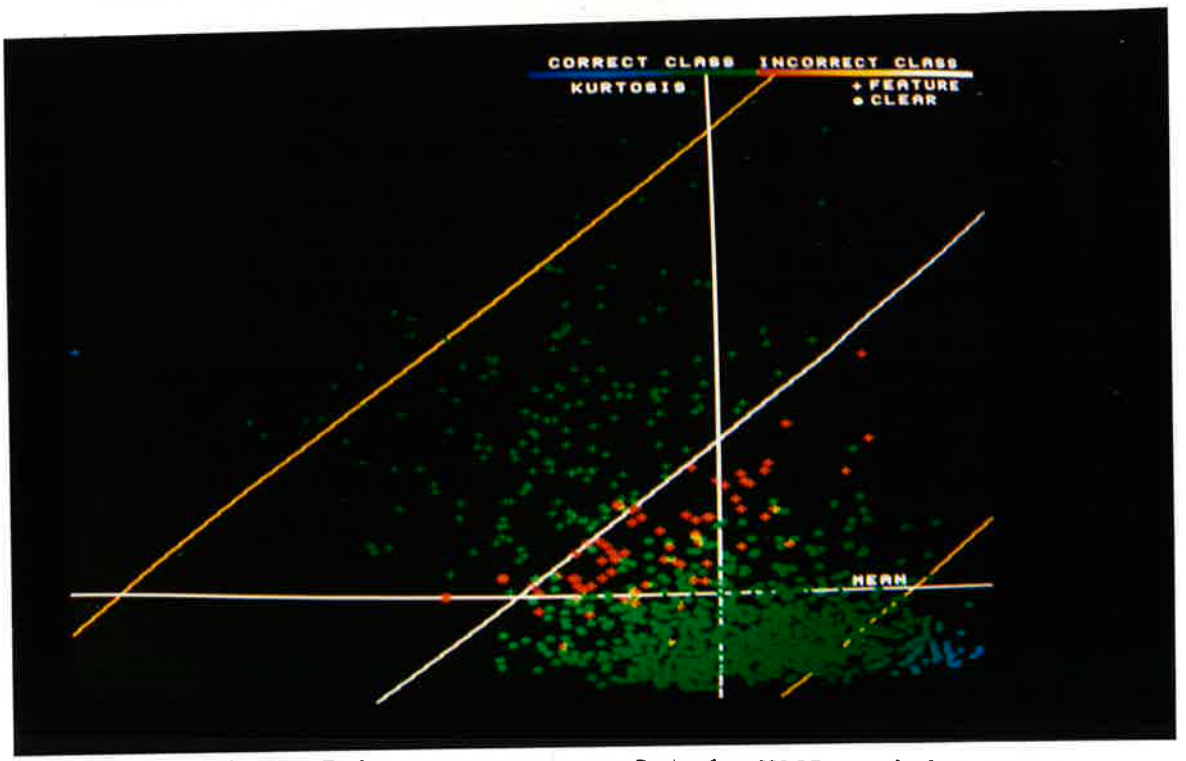

Figure 7.4 Pattern space for the WOOD training set.
Iteration = 0. **w** = (100.0, -100.0, 5.0)
73 samples misclassified: 70 false neg., 3 false pos.

Figure 7.4 shows the training set of the WOOD library of images where each sample point represents the mean and kurtosis feature measures of a local area of size 64 * 64 pixels. Each feature measure has been normalised using the mean and variance of the sample set and scaled to fit all the samples on the display screen. The axes are positioned at the mean value of each measure. The decision boundary itself can be drawn as a one dimensional line only if the parameter space is two dimensional. In this case it is drawn in the figures as a white line. The dead zone positions are parallel to the decision boundary and are

shown as yellow lines. In this way it is possible to see which samples influence the iteration of the cost function.

The colour scale indicates the relative distance of each sample from the hyperplane that separates the feature measure space into the two sections of the dichotomy. The two classes are represented by circles for clear local areas and crosses for feature local areas. Samples that are correctly classified are blue turning to green as they near the decision boundary. Samples that are misclassified by the decision boundary are coloured red and turn yellow to white as they move further from the boundary. This enables the status of any sample to be clearly perceived. Red circles are clear areas that are wrongly labelled as feature areas. They are called false positives. Red crosses are feature areas that are wrongly labelled as clear areas. They are called false negatives.

The iteration of the cost function can be visualised by inspecting the location of the decision boundary as the iteration proceeds. This is done using the two tonal measures of mean and kurtosis where each sample point is a vector in pattern space of the form:

$\mathbf{x}$ = (mean, kurtosis, 1)

Figure 7.4 shows the initial weight vector that represents the decision boundary. This is:

w_0 = (100, -100, 5.0)

The augmented coefficient, w_3 = 5.0, is the normal distance from the origin to the decision boundary. In this example the dead zone is dz = 10.0 which can be verified by comparing the distance from the origin to the decision boundary with the distance from the dead zone to the decision boundary. The latter is twice the former.

The parameters used in the iteration are:

dz = 10.0
lfn = 5.0
lossrat = 2.0
step = 0.002

The position of the decision boundary and the dead zone are shown in the following figures at intervals of 200 iterations of the algorithm up to 1000 iterations in figure 7.10. An interesting phenomenon is that the effective size of the dead zone decreases as the iteration proceeds. This is due to the increase in the value of the decision boundary (weight vector) coefficients and leads to the situation that less of the samples contribute to the iteration as it proceeds.

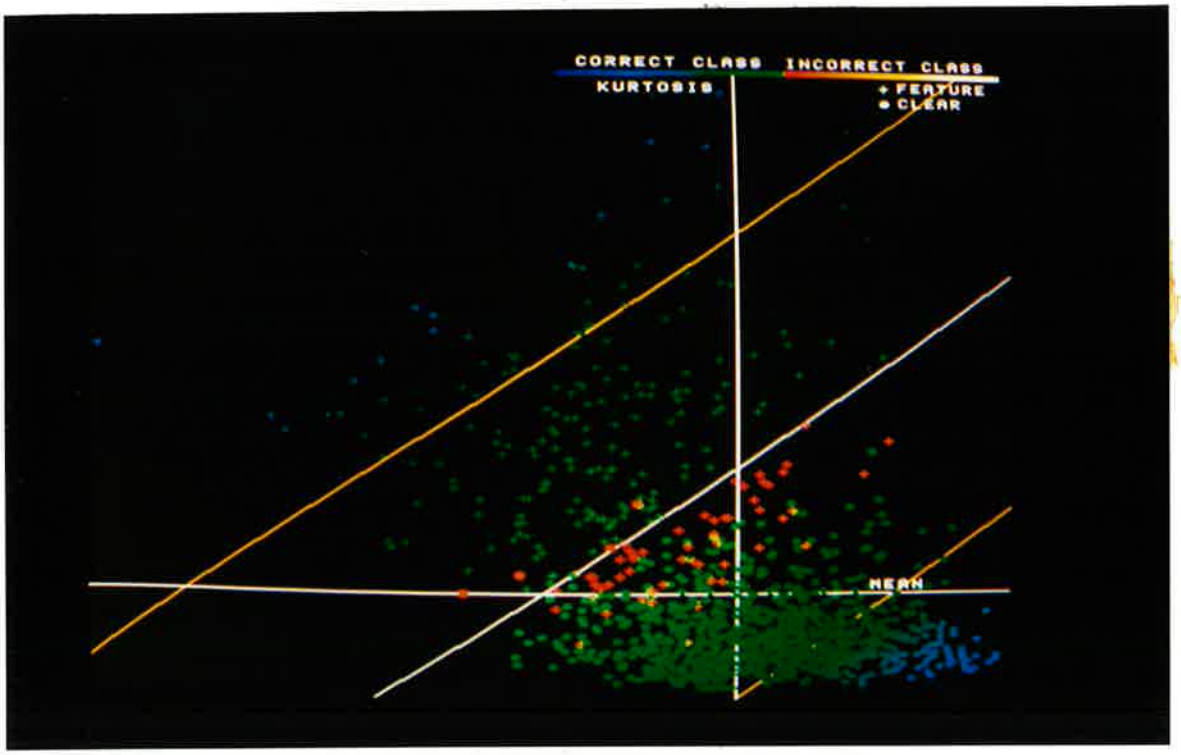

Figure 7.5 Pattern space for the WOOD training set.
Iteration = 200. **w** = (113.0, -132.1, 5.43)
64 samples misclassified: 60 false neg., 4 false pos.

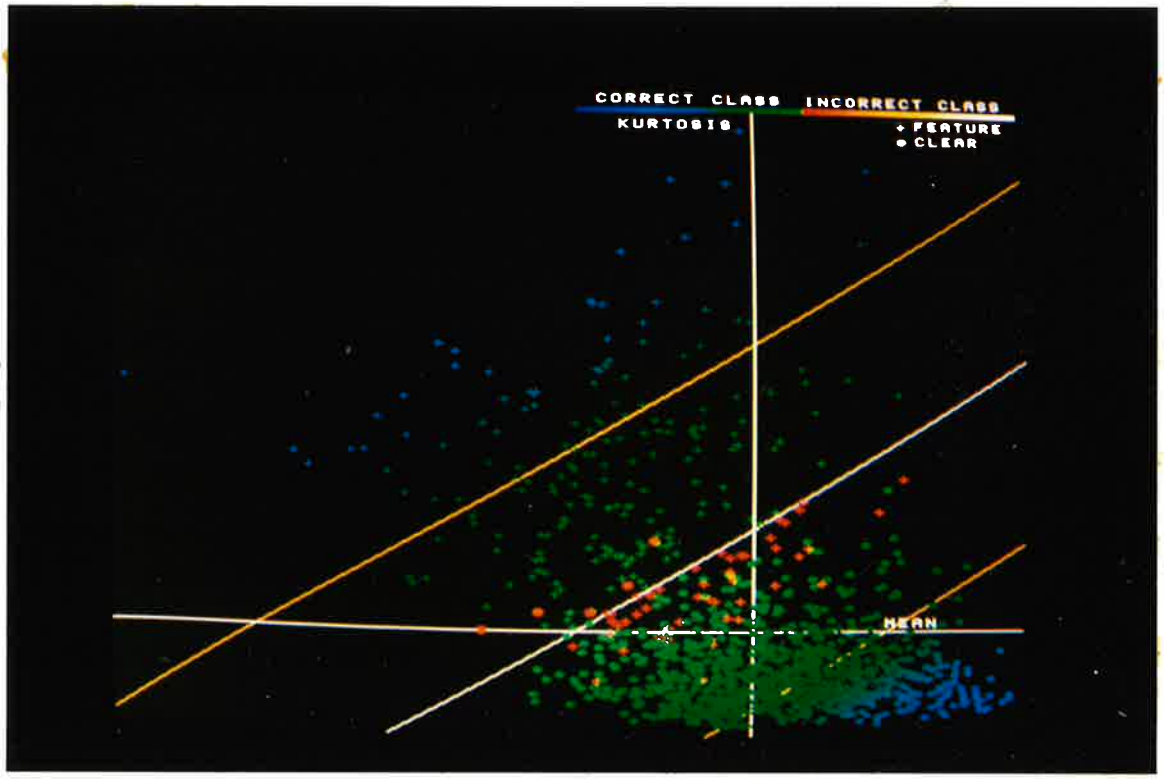

Figure 7.6 Pattern space for the WOOD training set.
Iteration = 400. **w** = (126.3, -168.1, 5.62)
57 samples misclassified: 52 false neg., 5 false pos.

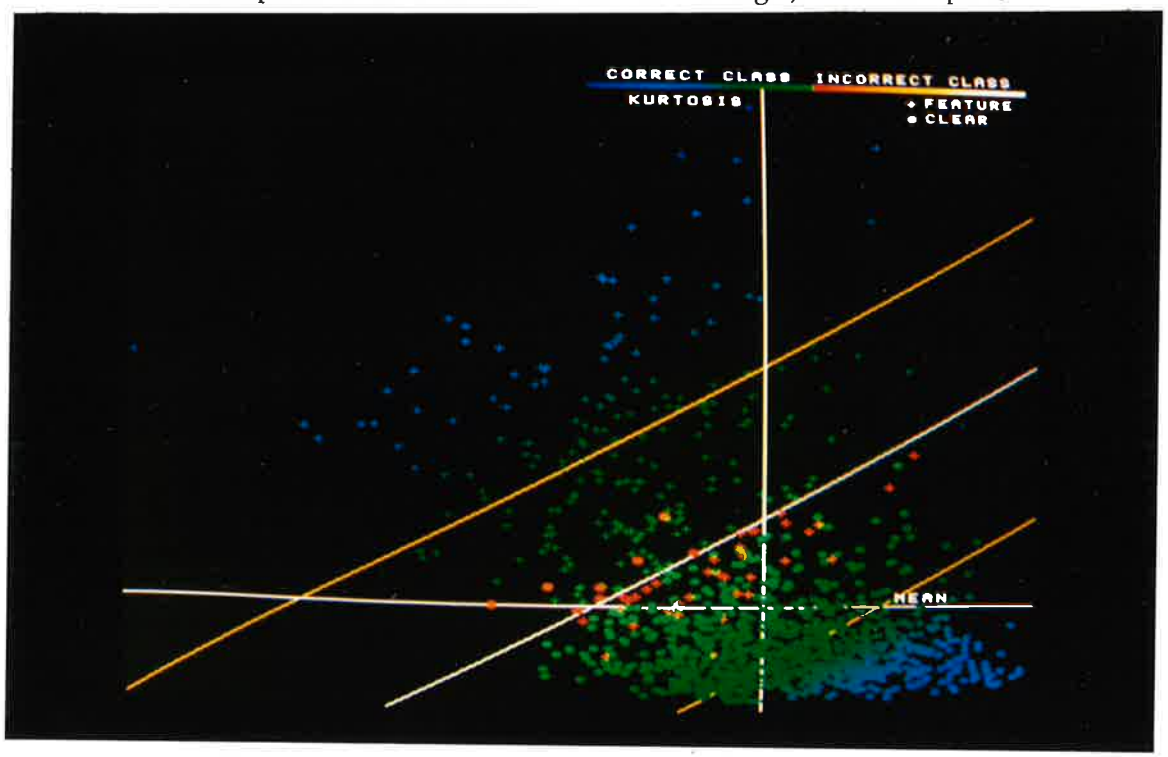

Figure 7.7 Pattern space for the WOOD training set.
Iteration = 600. **w** = (138.3, -203.3, 5.86)
48 samples misclassified: 40 false neg., 8 false pos.

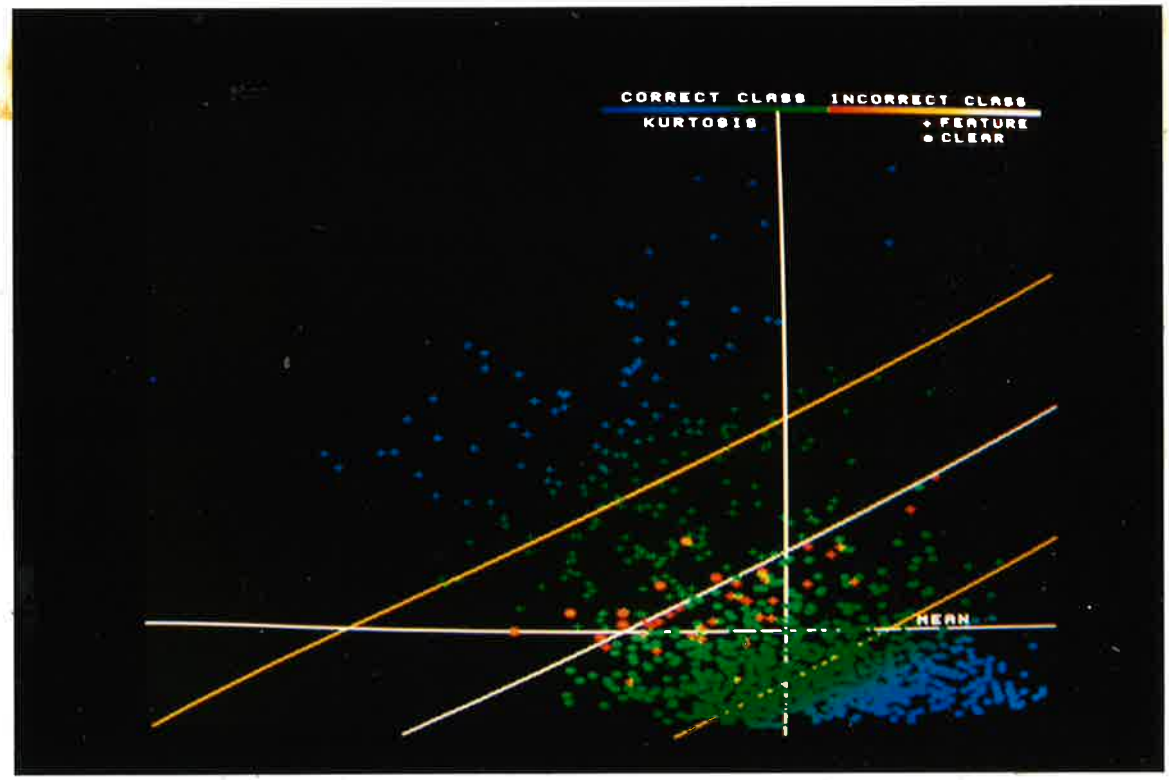

Figure 7.8 Pattern space for the WOOD training set.
Iteration = 800. w = (148.3, -233.1, 6.07)
44 samples misclassified: 36 false neg., 8 false pos.

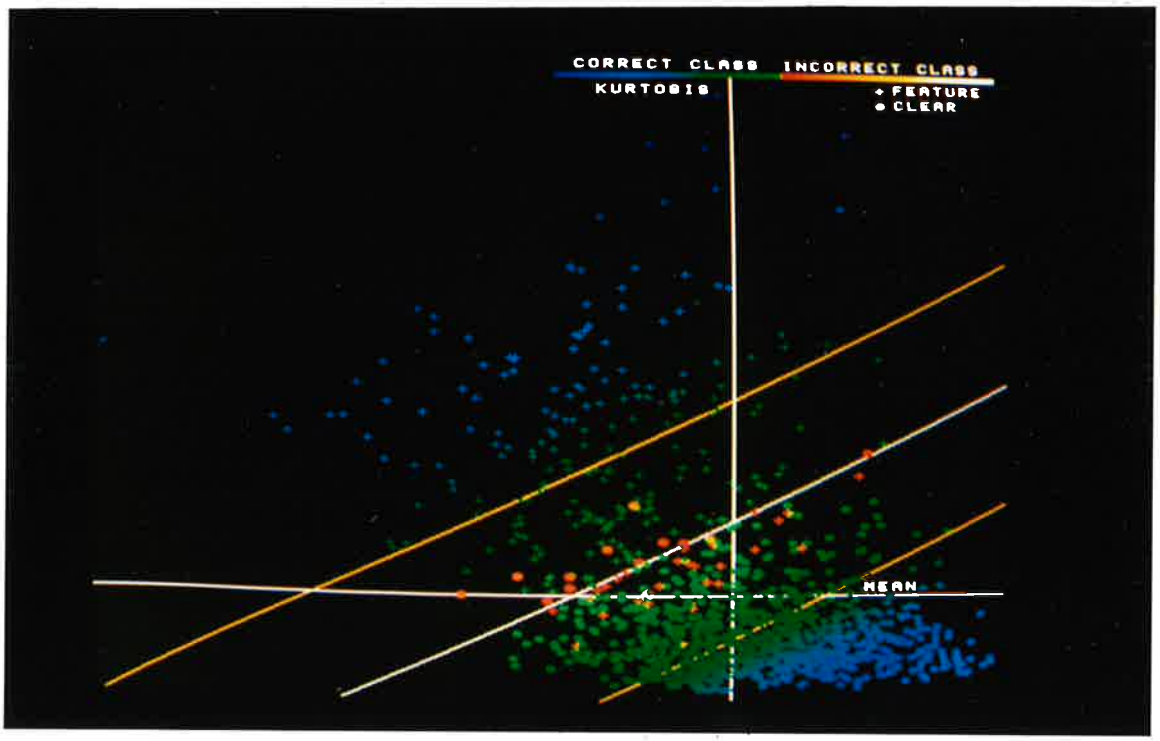

Figure 7.9 Pattern space for the WOOD training set.
Iteration = 1000. w = (156.1, -256.3, 6.25)
45 samples misclassified: 34 false neg., 11 false pos.

Figure 7.10 shows the effect of moving the decision boundary in a normal direction away from the origin by changing the augmented vector from 6.25 to 10.25. The number of false negatives is increased to 68 while the number of false positives is reduced to 2. The total number of samples misclassified is increased from 45 to 70. Figure 7.11 shows the effect of moving the decision boundary in the opposite direction. In this case the number of false negatives is reduced to 9 while the number of false positives is increased to 78. The total number of samples misclassified is increased to 87. This is a relatively easy means of changing the ratio of false positives to false negatives without calculating a new decision boundary by iteration. It is useful if the relative loss associated with each type of error is considered to be different from the assumption used in the iteration.

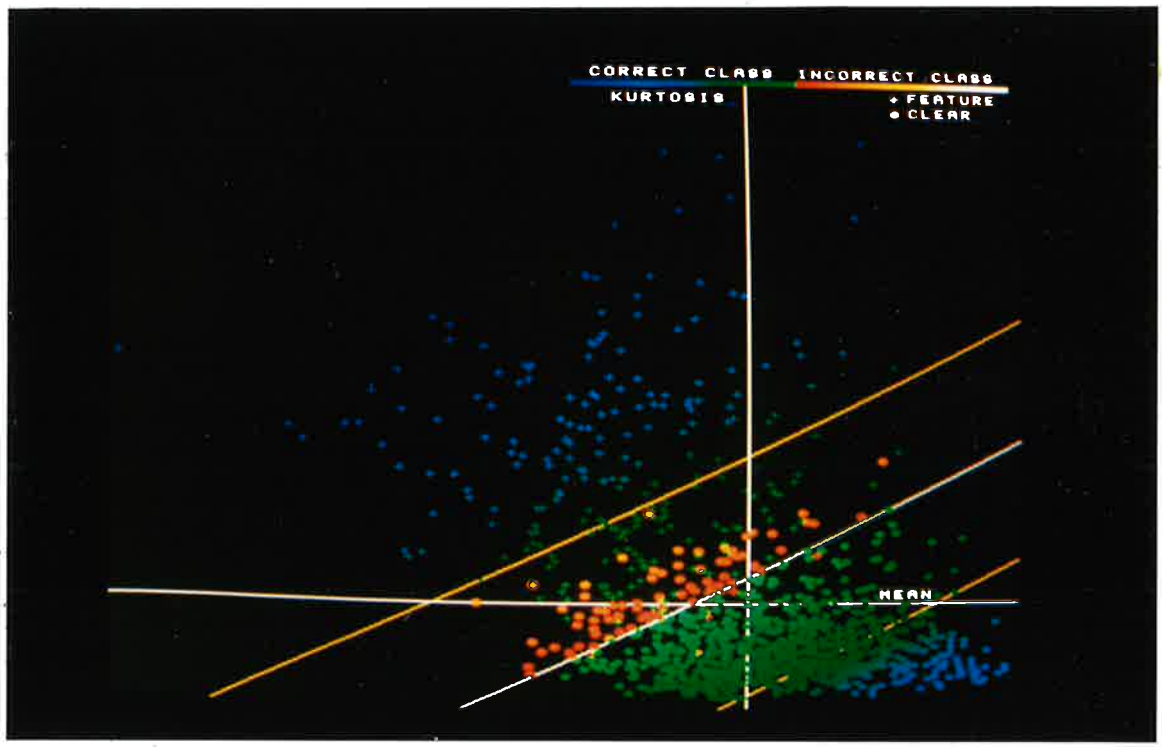

Figure 7.10 Pattern space for the WOOD training set.
w = (156.1, -256.3, 10.25)
70 samples misclassified: 68 false neg., 2 false pos.

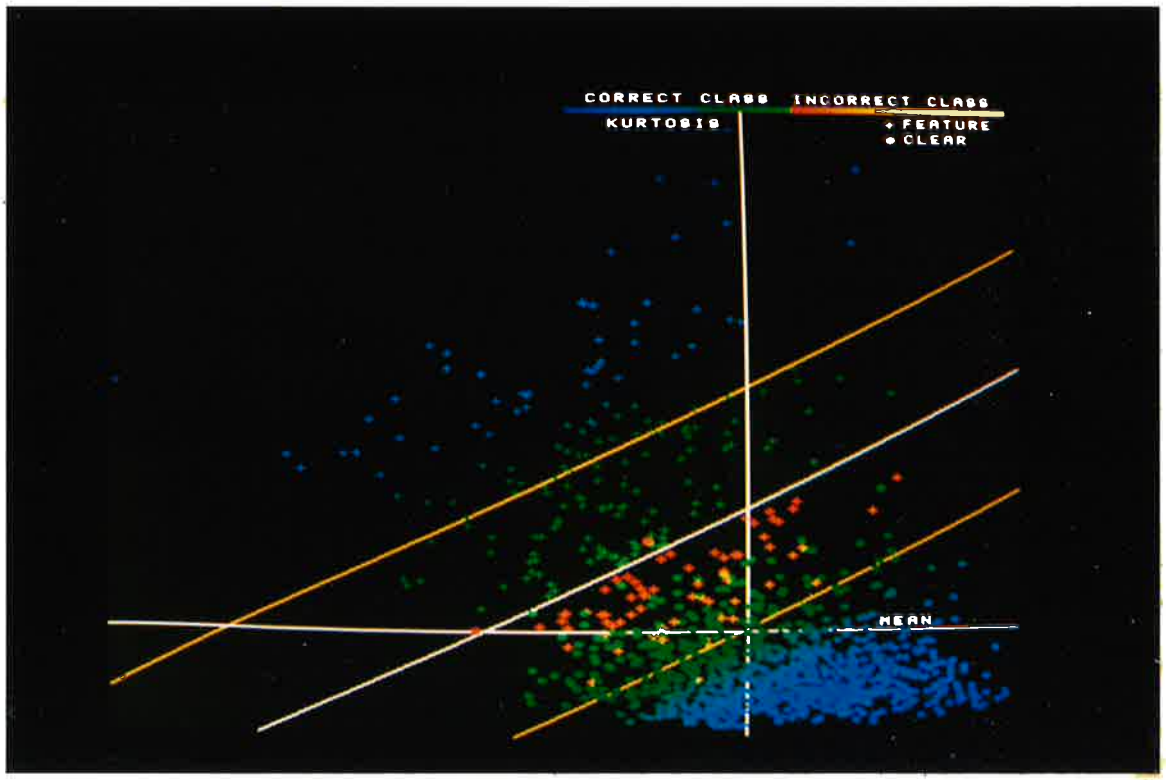

Figure 7.11 Pattern space for the WOOD training set.
w = (156.1, -256.3, 2.25)
87 samples misclassified: 9 false neg., 78 false pos.

Now that one is familiar with the position of samples in 2-d pattern space it remains to increase the dimension to 3-d space. Figure 7.12 shows the mean plotted against the standard deviation. This is plotted using the decision boundary derived for the previous figure using only the mean and kurtosis. The standard deviation is not included in the derivation of this decision boundary nor the calculation of the distance of each sample from the hyperplane. The samples are identically coloured in figure 7.12 and figure 7.9. The kurtosis axis can be considered to lie perpendicular to the page and the decision boundary rises out of the page at an angle dividing the samples into the two classes. By careful examination one can locate corresponding samples between the two figures by locating them along the common mean axis.

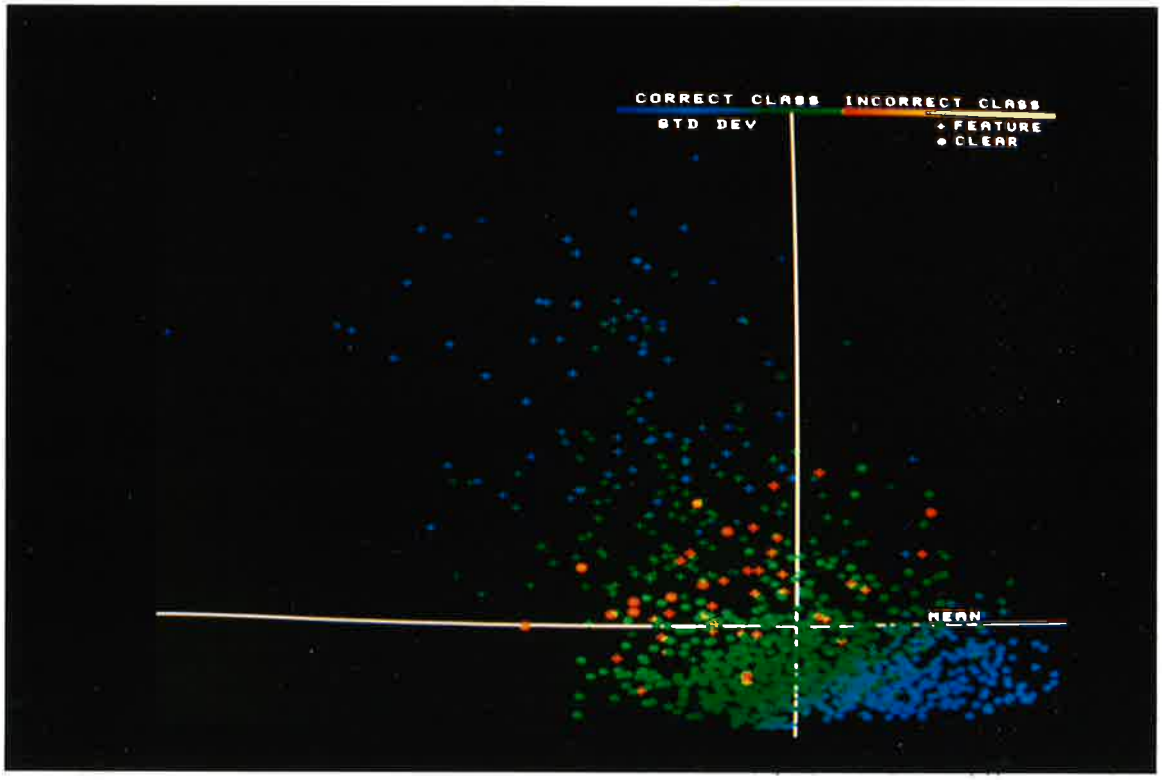

Figure 7.12 Pattern space for the WOOD training set.
Mean vs. Standard Deviation.
45 samples misclassified: 34 false neg., 11 false pos.

The extension into 4-d space is a little harder to conceptualise but it is able to be viewed in two dimensions at a time. This allows one to visualise the correlation or randomness of feature measures with each other in an overall sense.

8. INTRODUCTION OF TEXTURE MEASURES AND THE INITIAL ASSESSMENT OF THE CLASSIFICATION ALGORITHM

The pattern classification algorithm determines, by iteration, a decision boundary in multidimensional pattern, or feature measure, space. It reaches a decision boundary by trying to minimise a cost function that is applied to a known, labelled training set of samples. The feature measures of a new unlabelled sample position it in pattern space on one side or the other of the decision boundary, enabling it to be assigned a class.

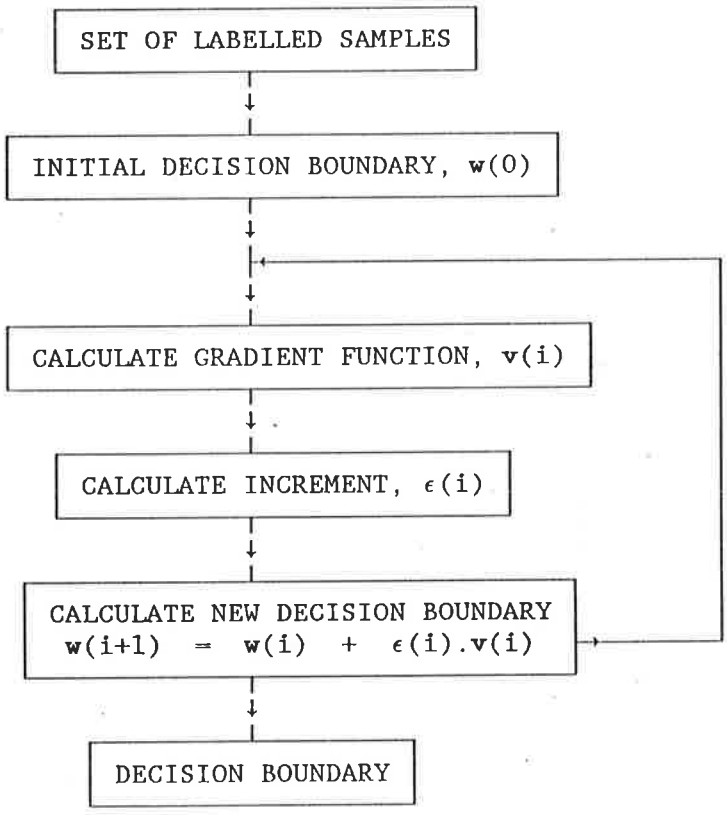

Figure 8.1 Flow diagram of the classification algorithm.

In order to calculate the position of the decision boundary in pattern space it is necessary to minimise a cost function to reduce the loss associated with misclassified samples of the training set. The minimisation of the cost function is an iterative procedure (figure 8.1) that is specified by a number of parameters. These are described in mathematical detail in the previous chapter and they are described in a less formal sense here with reference to the images of timber and the sample sets obtained from them.

The **dead zone (dz)** defines how many of the samples will contribute to the updating of the decision boundary. It is the normal distance to a pair of hyperplanes that are parallel to and either side of the decision boundary. All the samples that lie on the decision boundary side of their respective **dead zone** hyperplane contribute to the cost function iteration.

The **loss function parameter** (lfn) defines the loss associated with each misclassified sample as a function of its distance from the dead zone. The loss function defines that samples further from the dead zone contribute a greater amount to the iteration. A square function is used with the value of lfn defining the point at which the square function becomes a linear function (and so the slope of this linear function).

The speed of the iteration is determined by the **step**. The **step** is the proportion of the increment, which is calculated from the gradient function which, in turn, is calculated from the loss function. The larger the **step** the more rapidly the minimum of the cost function is approached with each iteration.

The **loss ratio** is the ratio of the loss incurred for misclassifying a feature as a clear area to the contrary misclassification. A value greater than 1 implies that the loss associated with the misclassification of a feature area is that many times greater than the contrary situation. For the initial assessment the value of the loss ratio is maintained at unity, i.e, lossrat = 1.

The former three parameters are varied over a range to determine the stability and performance of the algorithm with the

training set of samples of the FEAT library. Both the many-at-a-time and the one-at-a-time algorithm are examined and their relative performance compared. Using the four tonal measures and adding each of the texture measure in turn, an investigation is conducted into the usefulness of the texture measures to aid in the detection of feature local areas.

8.1 THE SPATIAL GREY LEVEL DEPENDENCE MATRIX (SGLDM)

Previous work in the visual inspection of timber has made use of textural information in addition to tonal information (Conners (1983, 1984)). The most useful measure of texture is the Spatial Grey Level Dependence Matrix (SGLDM) and is described in various references (Conners (1980), Haralick (1973a, 1973b, 1983), Ho (1968), Siew (1988), Van Gool (1985), Weszka (1976)). In order to compare the results of previous work and to examine the effectiveness of texture measures these are calculated for the sample set.

The SGLDM is based on the estimation of the second order joint conditional probability density functions, $f(i,j|d,\theta)$. Each $f(i,j|d,\theta)$ is the probability of going from grey level i to grey level j given an intersample spacing, d, and the direction given by θ. The directions are taken to be θ = 0°, 45°, 90°, 135°, 180°, 225°, 270°, 315°. Each of these joint conditional probability density functions can be written as a matrix:

$$\Phi(d,\theta) = f(i,j|d,\theta).$$

The size of these matrices is determined by the maximum difference in grey levels which is the number of grey levels itself. A simple relationship exists between certain pairs of matrices, $\Phi(d,\theta)$. Let the transpose of $\Phi(d,\theta)$ be $\Phi^t(d,\theta)$, then it can be seen that for a given d and θ the probability of going from grey level i to j is the same as going from grey level j to i. This can be expressed as:

$$\Phi(d,0) = \Phi^t(d,180)$$
$$\Phi(d,45) = \Phi^t(d,225)$$
$$\Phi(d,90) = \Phi^t(d,270)$$
$$\Phi(d,135) = \Phi^t(d,315)$$

Thus, half of the matrices are redundant and contribute no extra knowledge of the texture. By ignoring the distinction between opposite directions and utilising symmetric probability matrices the spatial grey level dependence matrices can be defined:

$$
\begin{aligned}
S_0(d) &= 1/2\ (\Phi(d,0) + \Phi(d,180)) \\
&= 1/2\ (\Phi(d,0) + \Phi^t(d,0)) \\
S_{45}(d) &= 1/2\ (\Phi(d,45) + \Phi(d,225)) \\
&= 1/2\ (\Phi(d,45) + \Phi^t(d,45)) \\
S_{90}(d) &= 1/2\ (\Phi(d,90) + \Phi(d,270)) \\
&= 1/2\ (\Phi(d,90) + \Phi^t(d,90)) \\
S_{135}(d) &= 1/2\ (\Phi(d,135) + \Phi(d,315)) \\
&= 1/2\ (\Phi(d,135) + \Phi^t(d,135))
\end{aligned}
$$

These matrices can be understood with reference to figure 8.1 which shows the directions of neighbouring pixels.

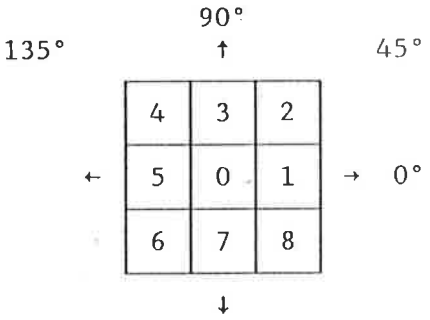

Figure 8.2 Spatial information parameter for intersample distance, d = 1. 0° neighbours to pixel 0 are pixels 1 and 5. 45° neighbours are pixels 2 and 6. 90° neighbours are pixels 3 and 7. 135° neighbours are pixels 4 and 8.

To illustrate the calculation of the SGLDM consider the 6 * 5 digital image of figure 8.3a. The image has 4 grey levels so the SGLDM's are 4 * 4 matrices.

0	0	1	1	1	3
2	2	3	3	2	2
1	1	2	1	3	1
3	0	3	0	1	1
0	3	1	3	1	0

Figure 8.3a A digital image of four grey levels.

For a value of d = 1 the summation SGLDM matrices for the four unique directions can be calculated:

$$P_0 = \begin{bmatrix} 2 & 3 & 0 & 4 \\ 3 & 8 & 2 & 6 \\ 0 & 2 & 4 & 2 \\ 4 & 6 & 2 & 2 \end{bmatrix} \qquad P_{45} = \begin{bmatrix} 2 & 1 & 2 & 1 \\ 1 & 4 & 3 & 6 \\ 2 & 3 & 0 & 3 \\ 1 & 6 & 3 & 2 \end{bmatrix}$$

$$P_{90} = \begin{bmatrix} 0 & 3 & 2 & 3 \\ 3 & 4 & 4 & 6 \\ 2 & 4 & 0 & 4 \\ 3 & 6 & 4 & 0 \end{bmatrix} \qquad P_{135} = \begin{bmatrix} 0 & 4 & 2 & 1 \\ 4 & 2 & 4 & 4 \\ 2 & 4 & 2 & 1 \\ 1 & 4 & 1 & 4 \end{bmatrix}$$

Figure 8.3b Summation matrices for the image in figure 8.3a. Appropriate normalisation is easily computed.

Normalisation of the matrices $P(d,\theta)$ give the probability density matrices $S_\theta(d)$. For an image of size (M,N) the normalistation factors are:

$\theta = 0$: (M - d) * N

$\theta = 45$: (M - d) * (N - d)

$\theta = 90$: M * (N - d)

$\theta = 135$: (M - d) * (N - d)

8.1.1 METHODS FOR REDUCING THE NUMBER OF GREY LEVELS

One can see that the size of the SGLDM matrices is dependent upon the number of grey levels in the image. If the original image is 512 * 512 pixels and 256 grey levels then the calculation of the SGLDM will require 4 matrices 256 * 256 of real numbers rather than the 8 bit integers of the image. This is clearly larger than the image itself. A method to reduce the number of grey levels without reducing the textural information in the image is described by Haralick (1973) and Conners (1978). Called **Equal Probability Quantising** (EPQ) it has been found to provide a near optimal way of reducing the number of grey levels in an image whilst retaining an accurate representation of the original image. It has the effect of normalising the image contrast.

EPQ is a method mapping the pixels in an image into an image with a smaller number of grey levels. The method is simply to place, as near as is possible, the same number of pixels in each of the new grey levels. It is not always possible to place exactly the same number in each of the reduced grey levels so a decision is made to place as near as possible to the same.

Let **X** be a non-negative random variable with cumulative probability function F_X. Let Q, the K-level equal-probability quantising function for **X**, be defined by $Q(x) = k$ if and only if $q_{k-1} \leq x < q_k$. We define q_0, q_1, q_2, ... , q_k in an iterative manner. Let $q_0 = 0$. Suppose q_{k-1} has been defined. Then let q_k be the smallest number such that:

$$\left| \frac{1 - F_X(q_{k-1})}{K - k + 1} + F_X(q_{k-1}) - F_X(q_k) \right|$$

$$\leq \left| \frac{1 - F_X(q_{k-1})}{K - k + 1} + F_X(q_{k-1}) - F(q) \right|$$

for all real q.

Ideally this will give a histogram of the reduced image with equal numbers of pixels in each of the grey levels. However,

that this may not be the most desirable form of data reduction because the majority of the "useful" information about wood features is in the lower end of the grey scale. Knots and bark are distinguishable by the dark portions of resin and bark while the large uniform areas of the background and growth rings have their information spread more or less evenly over the range of grey levels. An **unequal probability quantising** (UPQ) that places the dark areas of the original image over a greater range of grey levels in the reduced image than the lighter areas will tend to emphasise the darker areas.

To reduce the original image to $N_g = 16$ grey levels the following method is proposed. Given that $16 + 15 + \ldots + 1 = 136$ then the original image is mapped onto the the 16 grey levels such that the first grey level contains 1/136 of the image pixels; the second grey level contains 2/136 of the image pixels (or as close as possible); until the final sixteenth grey level contains 16/136 of the image pixels. The FEAT library of images are reduced to 16 grey levels in this manner for the purpose of calculating the SGLDM.

8.1.2 MEASURES OF TEXTURE

The following measures of texture are used in the pattern classification algorithm for the purpose of discriminating between clear and feature local areas.

Contrast:

$$CON(P_\theta(d)) = \sum_{i=0}^{N_g-1} \sum_{j=0}^{N_g-1} (i - j)^2 s_\theta(i,j|d)$$

Angular second moment:

$$ASM(P_\theta(d)) = \sum_{i=0}^{N_g-1} \sum_{j=0}^{N_g-1} [s_\theta(i,j|d)]^2$$

Correlation:

$$COR(P_\theta(d)) = \frac{\sum_{i=0}^{N_g-1} \sum_{j=0}^{N_g-1} (i - \mu_x)(j - \mu_y) s_\theta(i,j|d)}{\sigma_x \sigma_y}$$

Entropy:

$$ENT(P_\theta(d)) = \sum_{i=0}^{N_g-1} \sum_{j=0}^{N_g-1} s_\theta(i,j|d) \log s_\theta(i,j|d)$$

Local homogeniety:

$$LHOM(P_\theta(d)) = \sum_{i=0}^{N_g-1} \sum_{j=0}^{N_g-1} \frac{1}{1 + (i - j)^2} s_\theta(i,j|d)$$

Cluster shade:

$$CSH(P_\theta(d)) = \sum_{i=0}^{N_g-1} \sum_{j=0}^{N_g-1} (i + j - \mu_x - \mu_y)^3 s_\theta(i,j|d)$$

Cluster prominence:

$$CPR(P_\theta(d)) = \sum_{i=0}^{N_g-1} \sum_{j=0}^{N_g-1} (i + j - \mu_x - \mu_y)^4 s_\theta(i,j|d)$$

where:

$$\mu_x = \sum_{i=0}^{N_g-1} i \sum_{j=0}^{N_g-1} s_\theta(i,j|d)$$

$$\mu_y = \sum_{j=0}^{N_g-1} j \sum_{i=0}^{N_g-1} s_\theta(i,j|d)$$

and:

$$\sigma_x^2 = \sum_{i=0}^{N_g-1} (i - \mu_x)^2 \sum_{j=0}^{N_g-1} s_\theta(i,j|d)$$

$$\sigma_y^2 = \sum_{j=0}^{N_g-1} (j - \mu_x)^2 \sum_{i=0}^{N_g-1} s_\theta(i,j|d)$$

and where $s_\theta(i,j|d)$ is the (i,j)th element of $S_\theta(d)$ and N_g is the number of grey levels from which the SGLDM is extracted.

For a chosen inter-sampling distance, d, we have four angular spatial grey level dependence matrices. Hence we obtain four values for each of the seven measures described above. The mean of the four values are calculated as inputs to the pattern classification algorithm.

It is worthwhile to attempt to understand the significance of the texture measures, that is, what it is about the texture that each of them describe.

If the inter-sampling distance, d, otherwise known as the texture displacement vector, **tdv**, is of the order of the texture elements (texels) then the grey levels of points separated by tdv should often be quite different so that the values in $S_\theta(tdv)$ should be quite spread out from the main diagonal. Conversely, if the tdv is small and the texture is coarse then the difference between two pixels tdv apart will most

likely be small and the values in $S_\theta(tdv)$ will be clustered about the main diagonal. A measure of the scatter of values about the diagonal axis in $S_\theta(tdv)$ for various values of tdv will provide a measure of the texture coarseness. **Contrast (CON)** is the moment of inertia of the matrix S_θ about its main diagonal and so a large contrast corresponds to a spread out matrix and indicates a tdv that is of the order of the texel size. A low contrast corresponds to a texture coarser than the distance of the tdv.

The **angular second moment (ASM)** measure is smallest when all the values in the $S_\theta(tdv)$ matrices are as equal as possible. It is large when some matrix values are small and others large which occurs if they are clustered about the main diagonal or, alternatively, if they are clustered in any other part of the matrix. A low contrast ensures a high ASM but the converse of a high ASM does not neccesitate a low contrast if the clustering is other than the main diagonal.

The **entropy (ENT)** measure is largest for equal values in the $S_\theta(tdv)$ matrices which exists when the probability of going from one grey level to any other grey level is not very different. Such an image would have the appearance of noise with no regular structure. If the probability of going from one grey level to another is either very small or very large then a regular ordered structure is indicated and the entropy measure is small.

The **correlation (COR)** measures the degree to which the rows (or columns) of each matrix resemble one another. It is high when the values are uniformly distributed in the matrix and low if, for example, the values of the diagonal are small.

The **local homogeniety (LHOM)** measure weights most highly those matrix members that are on the diagonal and so gives a high value if the elements are clustered along the main diagonal. This corresponds to an image where there is little difference in grey levels between pixels spaced tdv apart. Widely varying pixel values will give a more scattered SGLDM and so a small local homogeniety measure. This measure can be used to determine the size of the texture elements.

The measures of **cluster shade (CSH)** and **cluster prominence (CPR)** are used by Conners (1983) in the study of timber inspection. They are included for comparative purposes. Cluster shade will give a large value if the values of the matrix are clustered along the diagonal at high i and j values, and will give low values if the clustering is on the diagonal at low i and j values. Cluster prominence is similar but will give high values if the clustering is at low or high i and j values.

8.2 APPLYING THE CLASSIFICATION ALGORITHM

A sample set is drawn from the FEAT set of images to form a training set from which to derive a decision boundary. A test set, in order to check the validity of the derived boundary, is to be drawn from the remaining images. The makeup of the training set is presented in table 8.1.

FEAT TRAINING SET
CONTAINS IMAGES FEAT01 TO FEAT32
512 LOCAL AREAS MADE UP OF
103 FEATURE AREAS
383 CLEAR AREAS
26 BACKGROUND AREAS

Table 8.1 Makeup of the FEAT sample set of local areas.

This training set comprises 10 large knots; 11 medium knots; 8 small knots; 14 pin knots; 7 pieces of resin or bark; 4 pieces of pith; and 1 example of needle trace. The images are divided into 64 * 64 pixel local areas and each area is given a classification. If the area contains a portion of a feature, no matter how small, it is given the classification of a feature local area. Clear areas of wood that contain only growth rings are given the classification of a clear local area. The background areas also comprise 8 fully uniform blank areas and 18

areas that are partially blank and partially contain wood. They arise due to the wood boards being different widths and some are not wide enough to fill the field of view of the camera. They are included to provide examples of such a situation occuring in a real time industrial situation and are placed in the clear area class.

Each image is filtered using a 3 * 3 low pass filter to eliminate a high frequency noise component that is inherent in the hardware. Experience has shown that this noise is most evident when the object is stationary, as in this case. If the object is moving the integrative nature of the charge coupled device sensing element will act in such a manner to provide a 'smooth' image. In this sense the software smoothing performed simulates the effect of capturing images of wood moving on a conveyor belt at speed. Figure 8.4 describes the low pass filter used. This filter leaves a border around the edge of the image of unprocessed pixels of width one pixel because there is no information beyond the border to input to the mask. This border is eliminated by copying the row/column of pixels next to the border onto the border itself.

1	2	1
2	4	2
1	2	1

Figure 8.4 Low pass filter. For each pixel in the original image this mask is placed centrally on top and each pixel multiplied by the corresponding mask value and summed. This value is divided by 16 and becomes a pixel of the filtered image in the position of the centre of the mask.

The eleven feature measures (four tonal and seven texture) are calculated for each disjoint 64 * 64 pixel local area for each image to form the training set. Each measure is normalised based upon the maximum and minimum of the range and

scaled between the values of 0 and 490 to facilitate an easy method of display.

Contrary to the previously defined nomenclature the feature areas form the set S_0 and the clear areas, including the background areas, form the set S_1. This has no effect on the performance of the algorithm except for a reversal of the signs of the coefficients of the weight vector.

8.2.1 STABILITY OF THE ALGORITHM ITERATION

To investigate the stability of the classification algorithm an initial starting point is needed. This is achieved by varying the parameters in an interactive version of the classification software until a stable iterative process is reached. Each local area is represented by a vector, **x**, in 11-dimensional (12-dimensional augmented) pattern space and is of the form:

x = (mean,std dev,skew,kurt,CON,ASM,COR,ENT,LHOM,CSH,CPR,1)

The initial weight vector that is determined using the four tonal measures and seven texture measures is given below for the many-at-a-time algorithm:

w = (-0.9664, 1.2089, -0.6693, 0.8728, -1.544, 0.8597,
1.4661, 1.0866, -1.3507, -2.9369, 2.2286, 6.2792)

Using a value of tdv = 1, and using a number of values for the parameters of dz, lfn, and step, the algorithm is cycled through 100 iterations and its performance noted. The behaviour is classed into five catagories in increasing order of instability:

STABILITY CLASSIFICATION	
1	STABLE AND STEADY
2	STABLE - R SLOWLY DECREASING
3	STABLE - R SLOWLY INCREASING
4	LARGE OSCILLATIONS IN R
5	VERY UNSTABLE OSCILLATIONS

where the misclassification distance, R, is the average of the distances of the samples from the dead zones. This is different to the cost function, R', which does not contain the norm of the weight vector as a denominator. R does not include the loss ratio in the loss function but still gives a measure of the cost of the misclassified samples. The misclassification distance, R, is dependent upon the value of dz, lfn, and the coefficients of the weight vector and so care must be exercised in comparing the value of R from one set of conditions to another. The measure of the misclassification distance, R, is described below:

$$R(w) = \frac{1}{M} \sum_{j=1}^{M} f\left(\frac{dz - w \bullet x_j}{||\omega||} \right)$$

where

$$f(z) = \begin{cases} 0, & z \le 0 \quad \text{zero loss, outside dead zone.} \\ z^2, & 0 < z \le lfn \quad \text{square function,} \\ 2.lfn.z - lfn^2, & z > lfn \quad \text{linear function.} \end{cases}$$

Both the many-at-a-time and the one-at-a-time algorithms were examined with variations in the parameters:

dz = 0, 50, 100, 200.

lfn = 50, 100, 200.

step = 0.1, 0.01, 0.001

and the results are presented in table 8.2 and table 8.3.

STEP	DEAD ZONE, dz = 0	50	100	200	
0.1	5	4	1	1	LOSS FUNCTION
0.01	2	2	3	2	lfn = 50
0.001	2	3	3	3	

STEP	DEAD ZONE, dz = 0	50	100	200	
0.1	4	4	2	1	LOSS FUNCTION
0.01	2	2	1	2	lfn = 100
0.001	2	2	1	2	

STEP	DEAD ZONE, dz = 0	50	100	200	
0.1	5	4	2	1	LOSS FUNCTION
0.01	2	2	2	2	lfn = 200
0.001	2	2	2	2	

Table 8.2 Stability of the many-at-a-time algorithm.

STEP	DEAD ZONE, dz = 0	50	100	200	LOSS FUNCTION
0.1	4	2	1	2	lfn = 50
0.01	2	2	3	2	
0.001	2	3	3	3	

STEP	DEAD ZONE, dz = 0	50	100	200	LOSS FUNCTION
0.1	4	2	2	2	lfn = 100
0.01	2	2	1	2	
0.001	2	2	1	2	

STEP	DEAD ZONE, dz = 0	50	100	200	LOSS FUNCTION
0.1	2	2	2	2	lfn - 200
0.01	2	2	2	2	
0.001	2	2	2	2	

Table 8.3 Stability of the one-at-a-time algorithm.

The most desirable situation is a stability factor of 2, where the value of R is slowly and steadily decreasing indicating that the decision boundary is iterating toward progressively more optimal positions. This situation is more likely to be achieved with combinations of small step sizes, large dead zones and large loss function parameters.

The one-at-a-time algorithm appears to be stable over a larger range of parameter values that the many-at-a-time algorithm although this is bought at a cost in speed. The many-at-a-time algorithm iterates toward the more optimal decision boundary at a faster rate than the former as the figures in table

8.4 clearly show for parameter values where the two algorithms are both stable.

	DEAD ZONE, dz			
	0	50	100	200
ΔD % many-at-a-time	0.8	0.3	0.13	0.2
ΔD % one-at-a-time	0.3	0.08	0.03	0.1

Table 8.4 Measure of the speed of iteration between the two algorithms after 50 iterations. lfn = 200; step = 0.01.

This is found to apply for all values of parameters where the two algorithms are both stable. The causes of instability are thought to stem from too large a change in the updating of the iteration. This can be caused by insufficient samples contributing toward the iteration as is the case if a small dead zone is used. It can also be caused by using a large step value which tends to create the unstable situation where the decision boundary oscillates between classifying all the samples as first one class and then the other. The effect of a large step can be tempered by the use of a large dead zone and a large loss function parameter, lfn.

8.2.2 EFFECT OF TEXTURE MEASURES

To measure the influence of texture measures in the ability to distinguish the dichotomy between feature and clear local areas the following method is followed. The parameters of the iteration are dz = 200; lfn = 200; and step = 0.1. The iteration is conducted 150 times using the many-at-a-time algorithm. The same initial weight vector is used as the previous section on stability:

w = (-0.9664, 1.2089, -0.6693, 0.8728, -1.544, 0.8597,
1.4661, 1.0866, -1.3507, -2.9369, 2.2286, 6.2792)

Using only those coefficients of w that apply, the iteration is performed using only the four tonal measures. The iteration is then repeated, this time using the four tonal measures plus the texture measure of contrast. This is repeated using each of the texture measures in turn with the four tonal measures. This complete process is repeated using a range of texture displacement vectors, tdv = 1, 2, 3, 4, 6, 8, 12. The results are presented in table 8.5 with the number of misclassified samples and the misclassification distance, R, after 150 iterations.

It is noted from table 8.5 that certain texture measures aid the tonal measures in the discrimination of the dichotomy of the sample set. A smaller number of misclassified samples is obtained with the texture measures of entropy, local homogeniety and cluster shade. The lowest value of M = 42 samples misclassified is an improvement of 5/512 or 1% of the sample areas from 9.2% misclassified to 8.2% (91.8% correct classification). Another way of expressing this is as a reduction of 11% in the number of misclassified samples.

The measures of contrast and angular second moment appear to neither help nor hinder the classification. Correlation and cluster prominence actively hinder the classificaton with more samples misclassified than there are feature areas. This implies that the coefficient of w that is used is incorrect and that a greater number of iterations are required to move it to a more optimal position. One would expect that if the feature measure has no additional information to improve the classification its weight vector coefficient will tend to zero to reduce its effect. This is observed to happen to a degree in the 150 iterations used in this test although the rate at which it occurs is much too slow to be revealed fully without many more iterations.

TONAL MEASURES No. MISCLASSIFIED, M = 47
MISCLASSIFICATION DISTANCE, R = 84

	tdv	1	2	3	4	6	8	12
TONAL + CON	M	50	49	49	49	50	50	51
	R	815	370	373	375	375	327	327
TONAL + ASM	M	51	50	50	50	50	44	50
	R	295	280	282	286	282	231	289
TONAL + COR	M	132	134	134	134	134	134	134
	R	237	273	273	273	275	276	277
TONAL + ENT	M	46	43	44	43	43	42	45
	R	175	292	224	244	291	286	229
TONAL + LHOM	M	44	44	44	42	43	44	42
	R	162	184	181	170	170	166	182
TONAL + CSH	M	53	47	45	46	47	45	46
	R	450	143	315	166	323	243	713
TONAL + CPR	M	105	111	109	110	111	112	111
	R	193	227	226	226	225	224	223

Table 8.5 The influence of the texture displacement vector, tdv, and the seven texture measures on the FEAT training set. lfn = 200; dz = 200; step = 0.1; 150 iterations.

A non-zero weight vector coefficient for a texture measure that does not significantly alter M means that the measure correlates with the classification performed by the tonal measures. That is, it does not contribute any new information that is relevant to the classification but at the same time it is not devoid of information. In fact it must contain information that is relevant to the classification or the non-zero coefficient would increase the number of misclassifications. One way to observe the difference in information content supplied by the new measure is to examine the samples misclassified both with

and without the measure and look for differences that are suggested by the nature of the texture measure.

The effect of the texture displacement vector is not entirely clear. The contrast measure has a large R with tdv = 1 but it drops to less than half for other values. The number of samples misclassified does not change significantly so the change in R must be due to the distribution of the samples in pattern space. This change in distribution is between the extreme values of the samples because of the method of normalisation of the measure. A possible explanation is that the low pass filtering of the image will tend to make adjacent pixels have less difference in grey levels between them (recall that the filter is a 3 * 3 mask). This will tend to increase the clustering of pixels about the main diagonal and reduce the range of the contrast pattern dimension with small values of tdv. When this range is scaled into pattern space based upon the maximum and minimum values in the range the samples will tend to be more spread out and so will contribute a greater amount to the value of R. The structure of the distribution in pattern space is not different because there is no marked difference in M but the form is noted by the difference in R.

This change in form is also exhibited with tdv = 1 for entropy and cluster prominence where the value of R is significantly different than with other values of tdv.

The angular second moment measure exhibits a drop in the number of samples misclassified with a value of tdv = 8. The corresponding value of R also shows a marked drop indicating that significant textural information is present at this position.

8.3 INITIAL CONCLUSIONS

The use of texture measures is found to hold the promise of improving the classification of local areas into feature and clear areas when combined with tonal measures. A reduction of 11% in the number of misclassified samples is possible using the single additional measure of either entropy or local homogeniety. Further work is required to determine the true effect of the

measures of correlation and cluster prominence. The initial coefficients used in the iteration appear to be inadequate to resolve an optimal decision boundary in the small number of iterations of the algorithm.

8.4 POSTSCRIPT

Further work on this data was confounded by the theft of computer hardware in 1987. However, the work of this section has been repeated and extended using a different data base (the WOOD library of images), collected with a different camera using different pieces of radiata pine boards. This work appears in the following sections and the correlation of results is encouraging. The differences in method, for example, normalisation and EPQ, are stressed where appropriate.

9. USING THE CLASSIFICATION ALGORITHM TO DETECT FEATURE LOCAL AREAS USING TONAL AND TEXTURE MEASURES

The initial results of feature detection with a pattern classification algorithm presented in the previous chapter reveal that tonal measures are capable of correctly classifying 90.8% of the local areas. The use of texture measures appears to be able to improve this marginally at a computational expense. In this section a larger data base of images is used and the parameters of the classification algorithm are examined more rigorously. After a greater understanding of the performance of the classification algorithm is obtained it is used to examine the usefulness of the various tonal and texture feature measures. The aim is to discover which measures best describe the features in radiata pine for the purpose of detecting feature areas.

The work is a development from that in chapter 8 in that a MATROX PIP-1024 frame-grabbing board is used instead of an IMAGEACTION board. This follows from the theft of the initial equipment.

9.1 DETERMINING THE INFLUENCE OF ALGORITHM PARAMETERS USING TONAL MEASURES

The initial work aims to establish the stability of the classification algorithm for a range of parameter values. This is done using the tonal measures of a sample set of images. Once a suitable range of parameter values are established, the usefulness of the tonal measures are determined by performing the iteration with different combinations of tonal measures. The attributes of each measure are revealed by examining the local areas that are misclassified.

9.1.1 DESCRIPTION OF THE TRAINING AND TEST SETS

The sample set comprising the WOOD set of images is divided into a training set, to derive a decision boundary, and a test set, in order to check the validity of the derived boundary. The makeup of these sets is presented in tables 9.1 and 9.2.

WOOD TRAINING SET	WOOD TEST SET
CONTAINS WOOD IMAGES	CONTAINS WOOD IMAGES
No. 01-19, 21-22, 24,	No. 20, 23, 25, 27-30,
26, 31-32, 38.	33-37, 39-48.
1073 LOCAL AREAS	604 LOCAL AREAS
239 FEATURE AREAS	189 FEATURE AREAS
834 CLEAR AREAS	415 CLEAR AREAS

Table 9.1 Makeup of the WOOD sample set of local areas.

FEATURE	TRAINING SET	TEST SET
LARGE KNOTS AND HOLES	18	11
MEDIUM KNOTS AND HOLES	9	15
SMALL KNOTS AND HOLES	19	20
PIN KNOTS AND HOLES	15	9
PITH	2	4
RESIN AND BARK	7	7

Table 9.2 Distribution of features in the training and test sets of the WOOD library of images.

Each image is filtered with the low pass filter that was used on the previous set of FEAT images. The local area size is again 64 * 64 pixels but a change in scale of the digitised images means that the size of the local area on the wood is

different. The scaling factors were measured from three images of graph paper that were captured at the start, middle and end of the session during which the WOOD images were recorded. The measured scaling factors are 0.533 mm/pixel in the horizontal (x) direction and 0.376 mm/pixel in the vertical (y) direction. This corresponds to a local area size of 34.1 * 24.1 mm on the wood surface.

Each local area is given a classification of either a feature local area, if it contains any amount of a feature, or a clear local area, if it contains only background area. This classification is very difficult for some areas which contain growth rings associated with diffuse spike knots.

The measures in the previous section were normalised using the extreme values of the range of each measure and scaled from 0 to 500 to facilitate display. The measures calculated for the WOOD set of images are normalised using the mean and standard deviation of each measure. This is considered to be more optimal than the former method (Sebestyen (1962)) because the presence of a small number of extreme values has less influence on the Euclidean distance between samples in pattern space. One significant effect of the normalisation procedure is that the weight vector coefficients and the classification algorithm parameters are different to those in the previous chapter. For each sample vector of the form:

$\mathbf{x}$ = (mean, std dev, skewness, kurtosis),

the normalisation values are:

mean = (185.21, 15.32, -18.08, 27.16),

standard deviation = (823.85, 341.24, 3053.65, 638.85).

Only the many-at-a-time algorithm is investigated using this data set because the results of the previous chapter show that this algorithm reaches a solution faster than the one-at-a-time algorithm.

Stability of the algorithm is investigated initially using the four tonal measures. By varying the parameters of the classification algorithm interactively a stable combination of algorithm parameters and weight vector coefficients is found from

which to explore. The sample vectors in pattern space are of the form:

$\mathbf{x}$ = (mean, std dev, skewness, kurtosis, aug),

and the initial weight vector is given the value of:

$\mathbf{w}$ = (100.0, -100.0, 100.0, -100.0, 5.0).

This is chosen as a standard starting point for all the iterations of the classification algorithm performed in this chapter and is the result of a considerable degree of experiment with different values. It becomes clear in the following pages that if the initial weight vector is too different from the ideal value or the parameter values are outside certain bounds, obtaining a sensible iteration is either extremely slow or impossible. The initial weight vector was obtained with the parameters of the classification algorithm equal to:

dz	= 10.0	Dead zone,
lfn	= 5.0	Loss function parameter,
lossrat	= 2.0	Loss ratio,
step	= 0.002	Iteration step.

Using the classification algorithm as outlined in the flow diagram of figure 8.1 this is found to provide a stable iteration with the iteration step and the measure of misclassification distance both decreasing.

Three measurements are used to measure the performance of the iteration process. They are best explained with reference to figure 9.2a. The heading of each graph lists the measures that are used in the iteration. In figure 9.2a the four tonal measures of mean, standard deviation, skewness and kurtosis are used. The upper left of the figure lists the parameter values that are used in the iteration as well as the number of iterations performed.

The plot on the upper right describes the number of samples misclassified as the iteration proceeds. This is the number of samples on the wrong side of the decision boundary. The number of samples that contribute to the iteration is greater and is defined by the position of the dead zone a distance dz from

the decision boundary. It can be seen in figure 9.2a that iteration actually increases the number of samples misclassified after about 470 iterations. This is not a degradation of the result if the misclassification distance is decreasing. The strength of this algorithm is that it is not dependant upon the number of samples misclassified but is able to pass through a position of a higher number of misclassified samples to reach a more optimal solution.

The plot at the lower left describes the value of the increment as the iteration proceeds. The increment is the actual size of the step taken at each iteration of the cost function. This plot is included to demonstrate the stability of the algorithm. The larger the increment the faster the algorithm will approach an optimum position, but too large an increment will lead to oscillations with many different samples being misclassified at each iteration. The increment is a function of the gradient of the cost function at each position of the decision boundary and in figure 9.2a the increment can be seen to be decreasing. This indicates that the gradient of the cost function is decreasing and the iteration is progressing toward a minimum.

The misclassification distance, R, is a measure of the cost of the samples that contribute to the iteration. It is explained more rigorously in the previous chapter and is plotted in the lower right of figure 9.2a. The slope of the curve indicates how rapidly the iteration is approaching the optimal minimum position of the decision boundary.

9.1.2 STABILTY WITH CHANGES IN DEAD ZONE AND LOSS FUNCTION

Using the same initial vector, w, the iteration of the pattern classification algorithm is performed for a number of combinations of dz and lfn whilst maintaining the loss ratio and step constant. The dead zone (dz) is given values of 5.0, 10.0, 15.0, and 20.0. These are matched with values for the loss function parameter (lfn) of 5.0, 10.0, 15.0, and 20.0. The value

for lossrat = 2.0; and step = 0.002; and the iteration is performed 1000 times.

Figures 9.1(a-d) to 9.4(a-d) show the results obtained. Variation of the dead zone with lfn constant can be observed in, for example, figure 9.1a, 9.2a, 9.3a, and 9.4a. Similarly, variation of the lfn with a constant dead zone can be observed, for example, in figures 9.2a, 9.2b, 9.2c, and 9.2d.

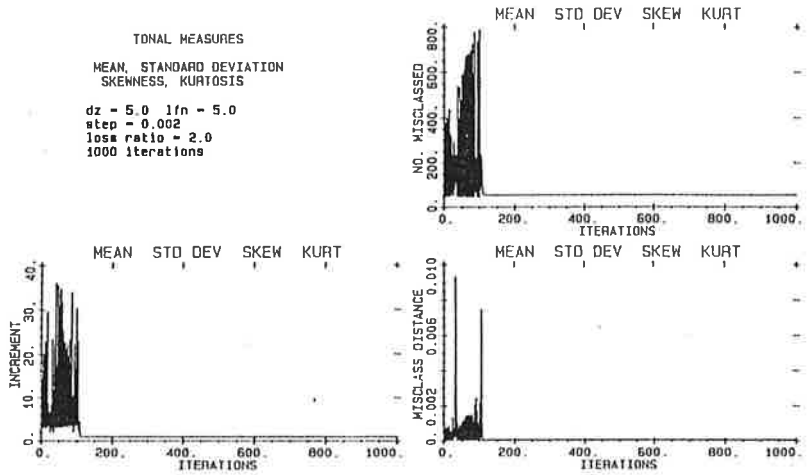

Figure 9.1a dz = 5.0; lfn = 5.0.

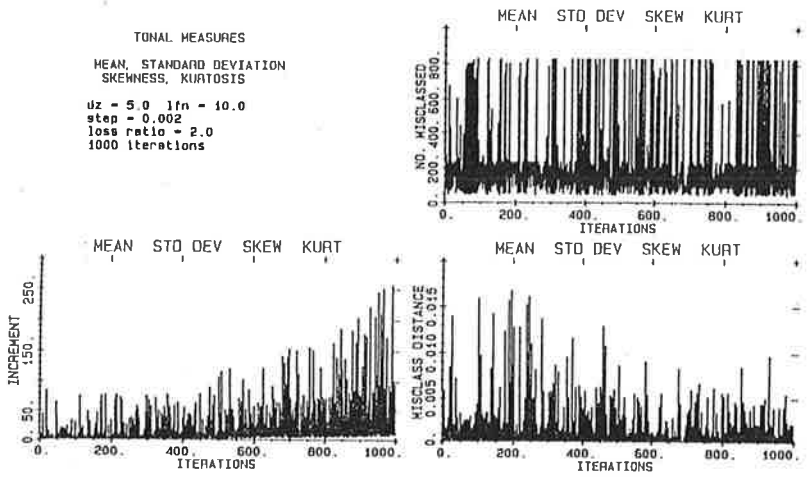

Figure 9.1b dz = 5.0; lfn = 10.0.

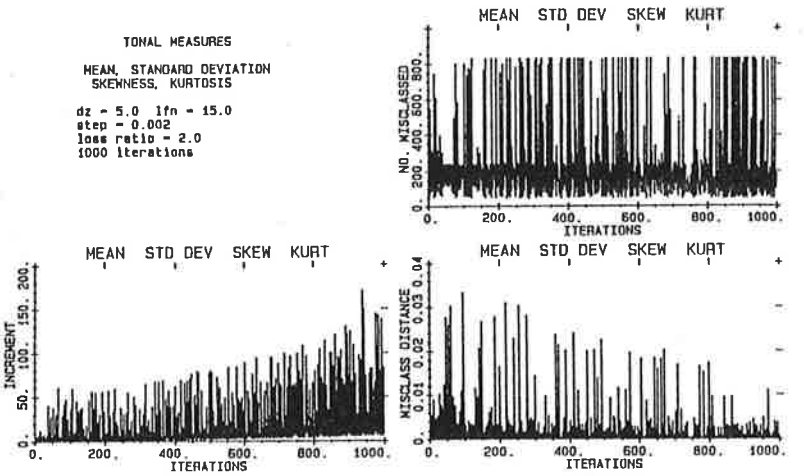

Figure 9.1c dz = 5.0; lfn = 15.0.

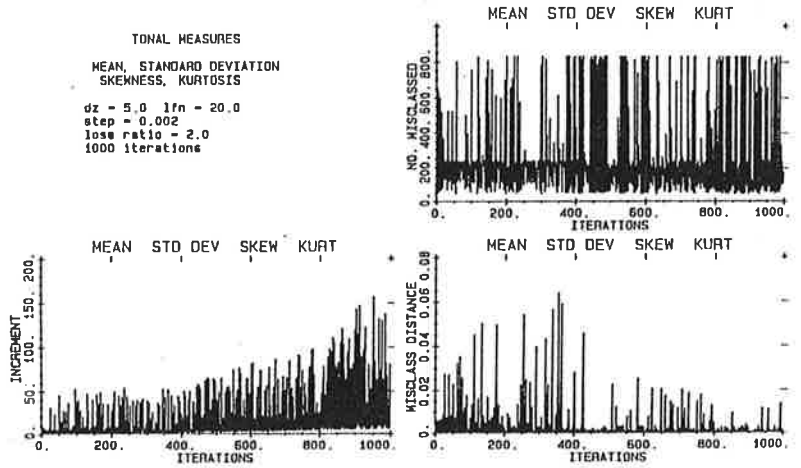

Figure 9.1d dz = 5.0; lfn = 20.0.

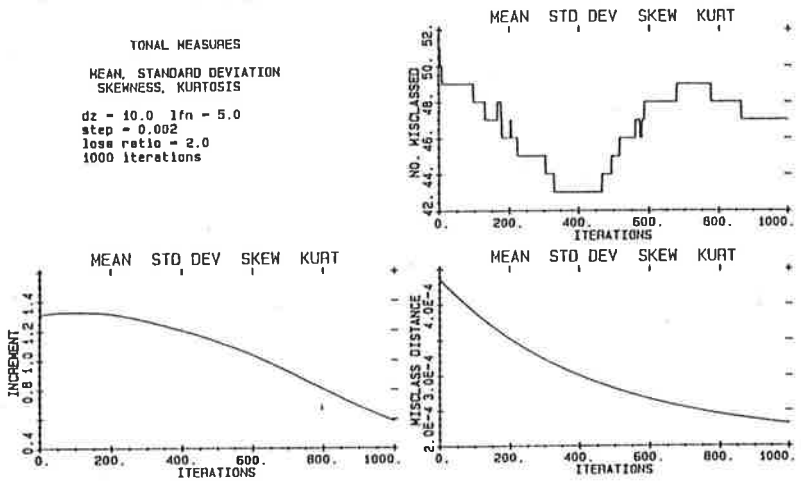

Figure 9.2a dz = 10.0; lfn = 5.0.

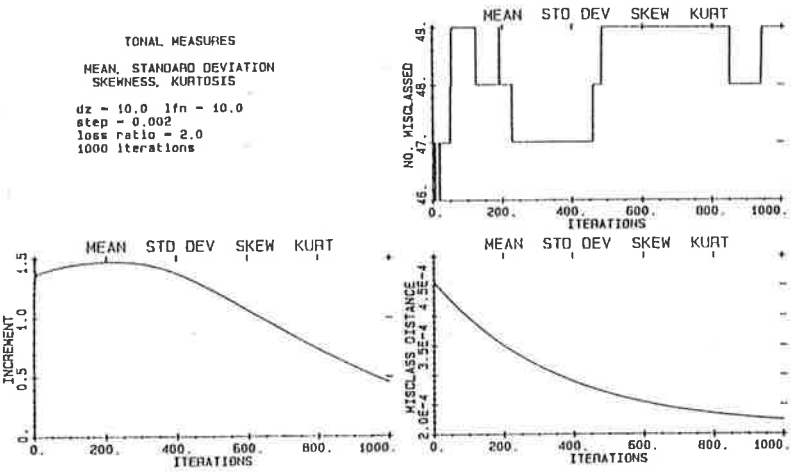

Figure 9.2b dz = 10.0; lfn = 10.0.

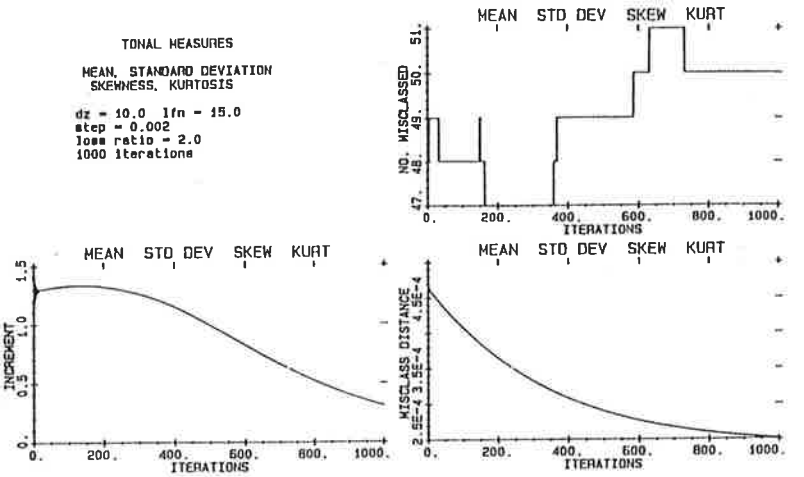

Figure 9.2c dz = 10.0; lfn = 15.0.

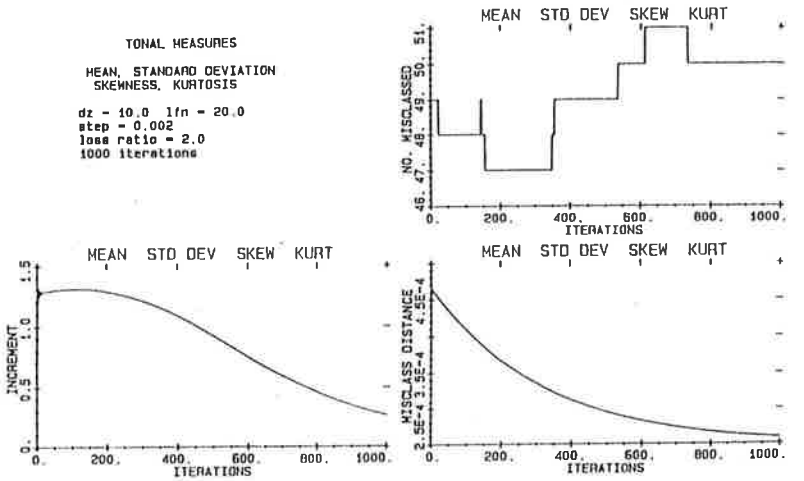

Figure 9.2d dz = 10.0; lfn = 20.0.

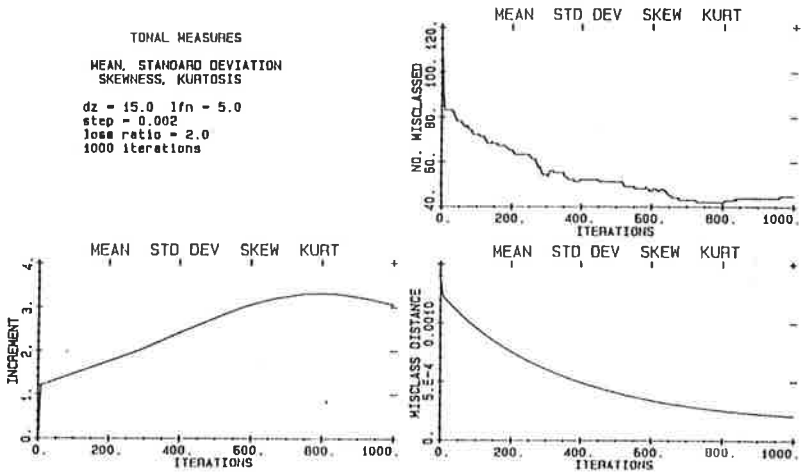

Figure 9.3a dz = 15.0; lfn = 5.0.

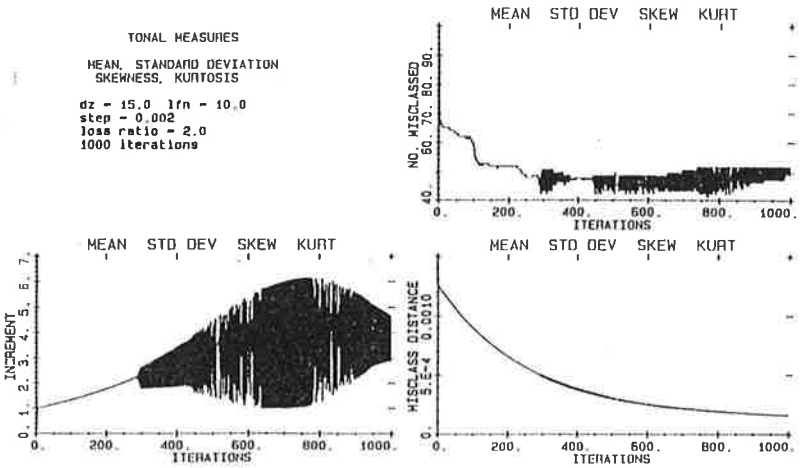

Figure 9.3b dz = 15.0; lfn = 10.0.

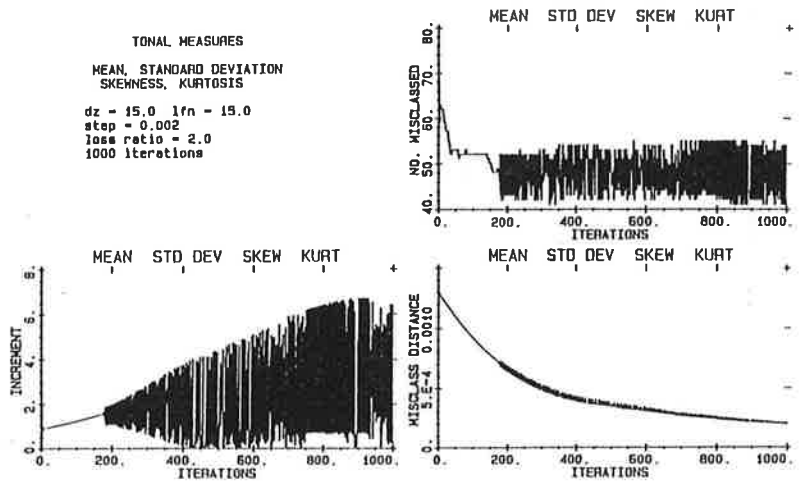

Figure 9.3c dz = 15.0; lfn = 15.0.

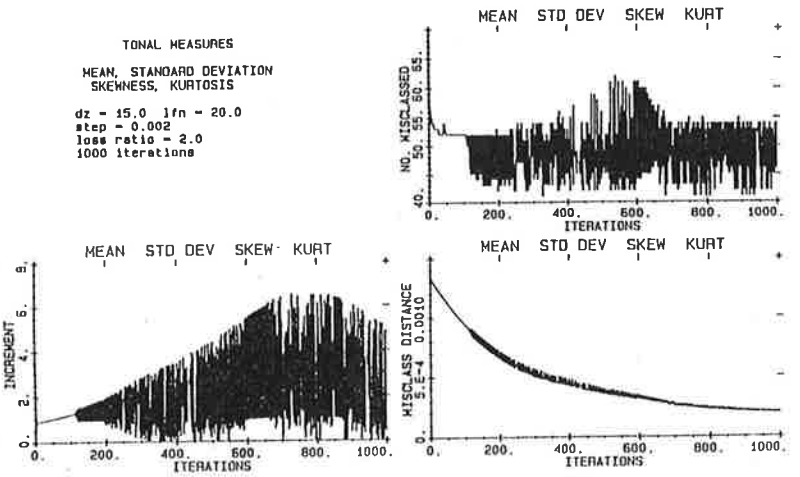

Figure 9.3d dz = 15.0; lfn = 20.0.

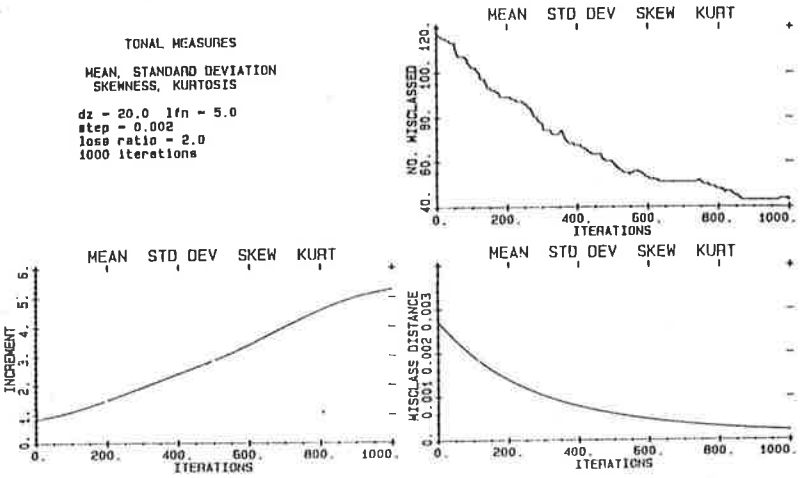

Figure 9.4a dz = 20.0; lfn = 5.0.

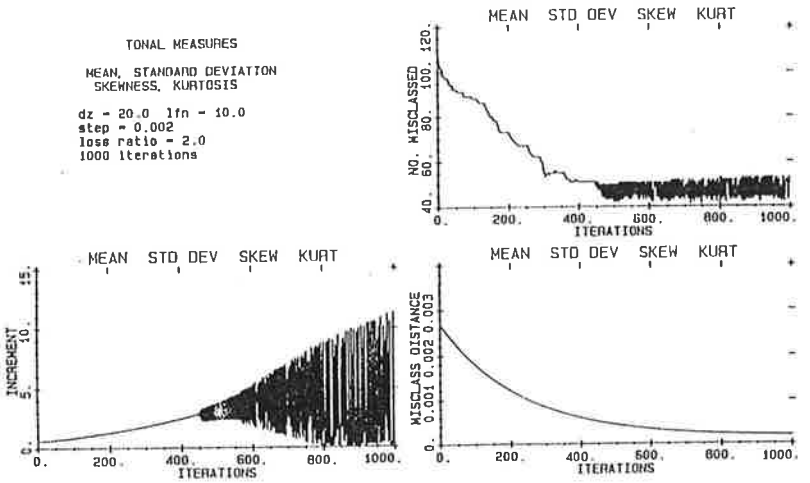

Figure 9.4b dz = 20.0; lfn = 10.0.

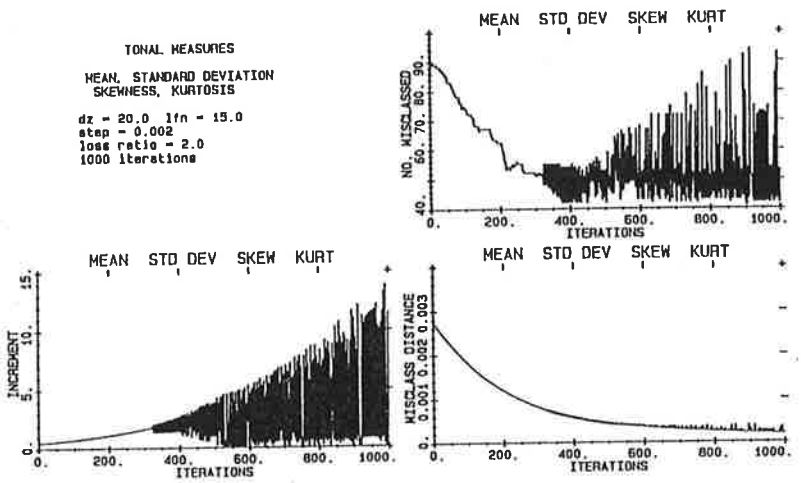

Figure 9.4c dz = 20.0; lfn = 15.0.

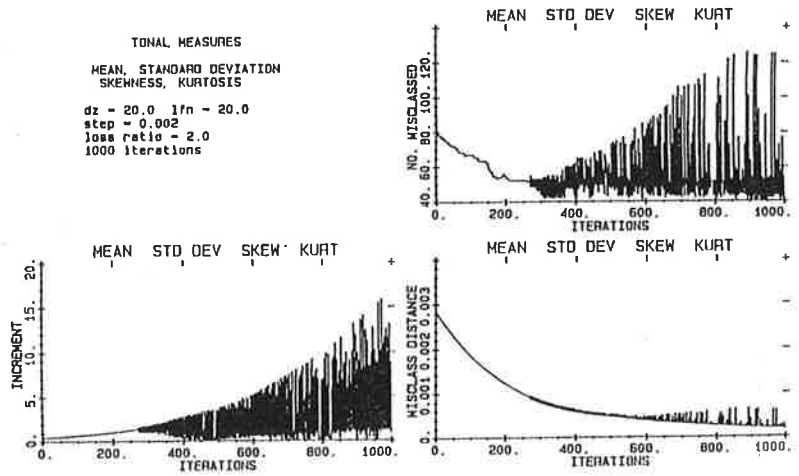

Figure 9.4d dz = 20.0; lfn = 20.0.

Oscillatory behaviour exhibits itself in a number of the plots and is usually evident in all three measurement parameters at once but different effects can be discerned.

The training set has 1073 samples made up of 834 clear areas and 239 feature areas. In the unstable graphs of figures 9.1(b-d) the oscillations are between 239 and 834 misclassified samples. This can be interpreted as the hyperplane having all the samples alternatively on one side and then the other. This unstable behaviour due to a small value of the dead zone implies that there are not enough samples contributing to the iteration. This gives rise to a large increment value which produces the unstable behaviour.

In figures 9.2(a-d) and 9.3(a-d) the increment shows a general trend to increase to a maximum and then decrease toward zero as the optimum position of the decision boundary is neared. Oscillations of the iteration, as in figures 9.3(b-d), tend to be modulated by this trend. Only the unstable oscillations of figures 9.1(b-d) do not conform to this generality, having increment values exceeding 20 and up to 200 with the misclassification distance being large and unstable as well.

Common points to be drawn are:

a) that oscillations are most likely to occur when the increment becomes greater than 1.5, and

b) if the increment value becomes too large (>10) the oscillations may degenerate into instablility.

The misclassification distance shows a general trend to decrease to a minimum asymptotically. Oscillations in the increment are usually coupled with oscillations in the misclassification distance curve and, just as the oscillations in the increment curve are modulated by a general trend, so too are the oscillations in the misclassification distance curve modulated by its asymptotically decreasing shape. This can be seen in figures 9.3(b-d).

Figures 9.1(b-d) show that the small value of the dead zone is unstable, with all three measurement parameters showing large, unstable oscillatory behaviour. Figure 9.1a shows an initial period of instability that resolves into stability. It appears that a small dead zone includes too few samples to be able to iterate efficiently.

Figures 9.3(b-d) and 9.4(b-d) exhibit instability after an initial stable period but the increment can be seen to be modulated suggesting that stability will be regained if the iterative process is allowed to continue. Figures 9.4c and 9.4d which combine a large dead zone with a large loss function parameter do not suggest that the increment is reaching a peak and unstable behaviour with further iterations would not be unexpected. The large and increasing oscillations in the number

of samples misclassed and the misclassification distance lend further evidence to this view.

Increasing the dead zone more than the value of 10.0 tends to lead to oscillatory and unstable behaviour unless the lfn is very small (5.0) reducing the influence of the samples. A large dead zone, which includes a greater proportion of the samples in the iteration, tends to reduce the ability of the algorithm to iterate effectively. Combined with a larger loss function parameter this ability is further reduced and the oscillations are larger. Increasing the loss function parameter increases the contribution of the samples furthest from the dead zone and increases the instability of the iteration, particularly with larger values of the dead zone. Increasing the dead zone increases the peak value of the increment and it has already been noted that if the increment is above 1.5 instability is likely.

The most stable plots are figures 9.2(a-d) with a dead zone of 10.0 and figures 9.3a and 9.4a with a loss function of 5.0, although the latter two show a peak increment value that is associated with instability. The parameters of figure 9.2a (dz = 10.0, lfn = 5.0) are chosen as the master values around which variations of the other parameters are explored. Dead zone values of 5.0 are unstable and values of 15.0 and over are also unstable. Values of the loss function parameter above 5.0 also tend to contribute to instability when combined with a dead zone greater than 10.0. The results suggest that the optimal values of dz = 10.0 and lfn = 5.0 will ensure a stable iteration with a reducing iteration increment and a minimum misclassification distance.

9.1.3 VARIATION OF THE LOSS RATIO

The reason the loss ratio is introduced is to be able to influence the position of the decision boundary with the perceived difference in loss associated with false negatives and false positives. In the timber inspection process it is felt that the loss incurred with a false negative (and so not detecting a

feature local area) is greater than the loss incurred with a false positive, which requires that further processing be done to confirm the local area as clear. Changing the value of the loss ratio should lead to a change in the proportion of false negatives to false positives. Reference to table 9.3 indicates that this is the case.

LOSS RATIO	MISCLASS DISTANCE ($*10^{-6}$)	NUMBER SAMPLES MISCLASSED	FALSE NEG	FALSE POS
1.0	98.6	42	32	10
1.5	430.8	60	55	5
2.0	231.3	47	29	18
2.5	187.1	56	27	29
3.0	169.1	59	23	36
3.5	161.3	58	18	40
4.0	157.7	61	16	45

Table 9.3 Variation of the ratio of false positives to false negatives over a range of values of the loss ratio. dz = 10.0; lfn = 5.0; step = 0.002; 1000 iterations; 1073 samples; 239 feature areas; 834 clear areas.

The behaviour of the algorithm with the variation of the loss ratio from 1.0 to 4.0 is shown in figures 9.5(a-g). Figure 9.5a shows the results with a loss function of 1.0. This means that the value of loss associated with a misclassified sample is the same for a false negative and a false positive. Increasing the loss ratio increases the loss associated with a false negative to a multiple of that for a false positive an equal distance from the dead zone. This has an effect similar to increasing the dead zone in that some of the misclassified samples contribute more to the iteration and so the increment is larger. This leads to the situation of figure 9.5g where the

increment is large enough for the algorithm to enter an oscillatory phase. This is not unstable however and it recovers to reach a minimum. Figure 9.5a enters an unstable period but it too recovers.

Figure 9.5b exhibits stable behaviour but rapidly reaches a position with a zero increment and no change which is unique among the plots studied. This behaviour is consistent with the presence of a local minimum which prevents the gradient function from incrementing toward the global minimum. It is found that this problem can be overcome by a suitable change in the coefficients of the initial weight vector.

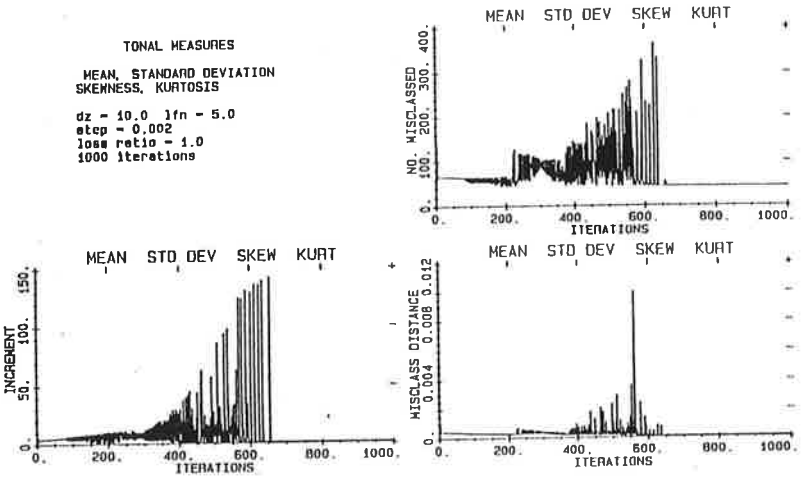

Figure 9.5a Loss ratio = 1.0.

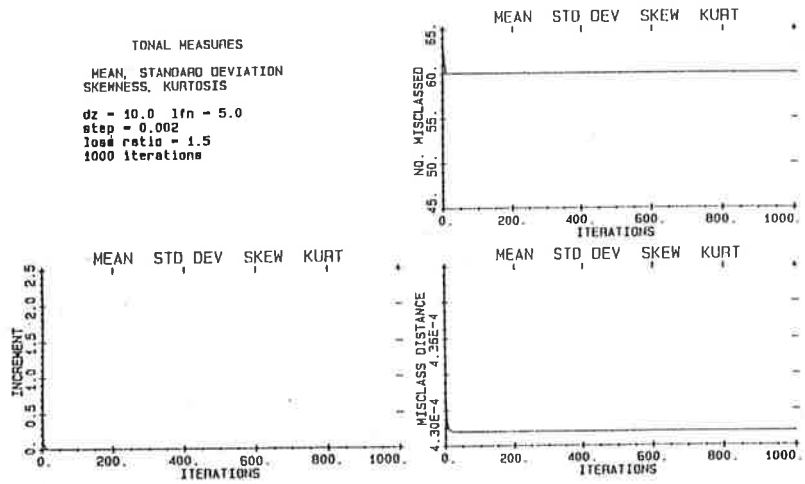

Figure 9.5b Loss ratio = 1.5.

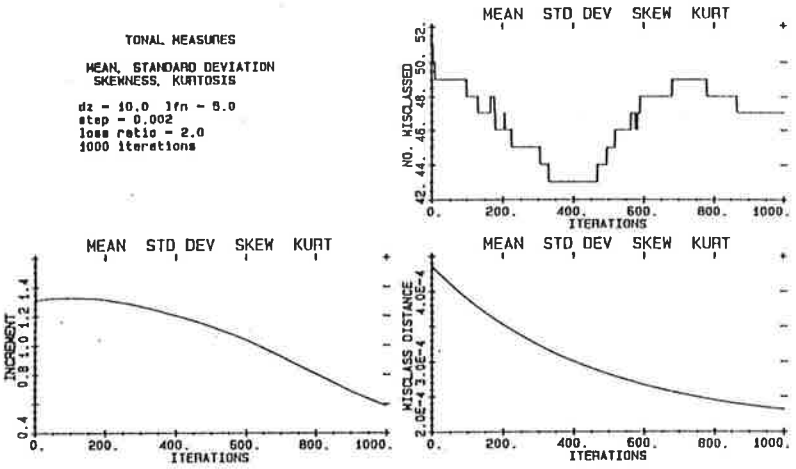

Figure 9.5c Loss ratio = 2.0.

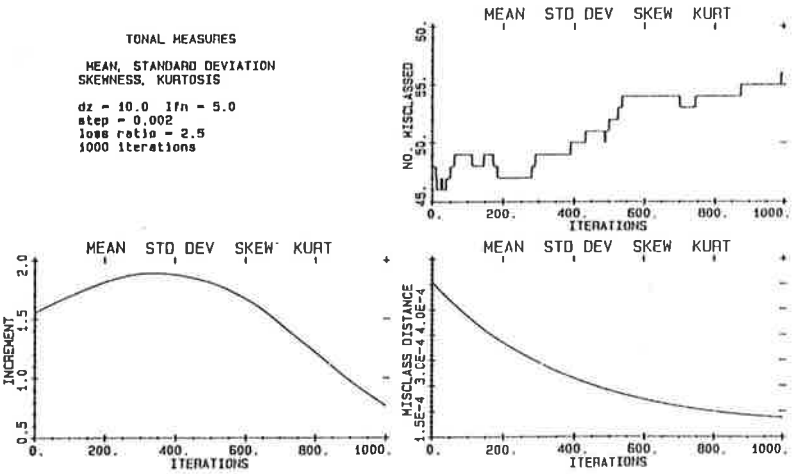

Figure 9.5d Loss ratio = 2.5.

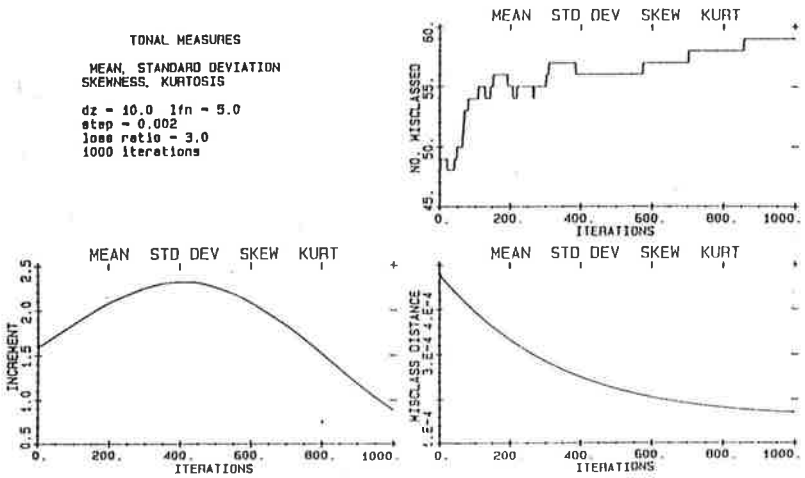

Figure 9.5e Loss ratio = 3.0.

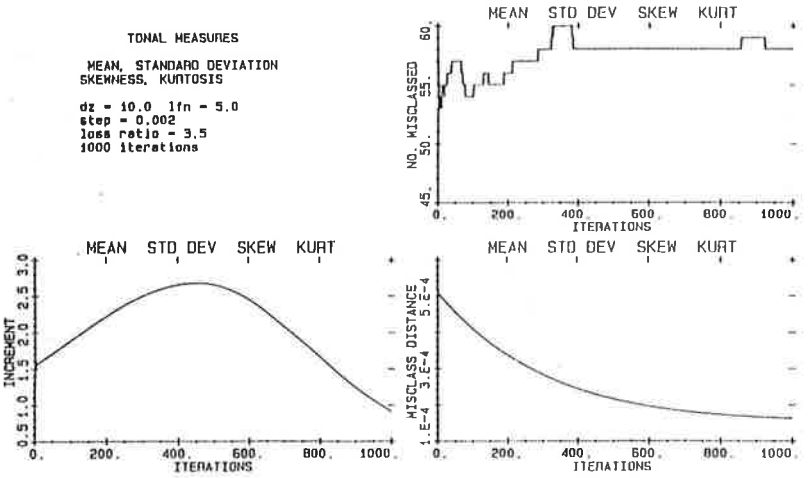

Figure 9.5f Loss ratio = 3.5.

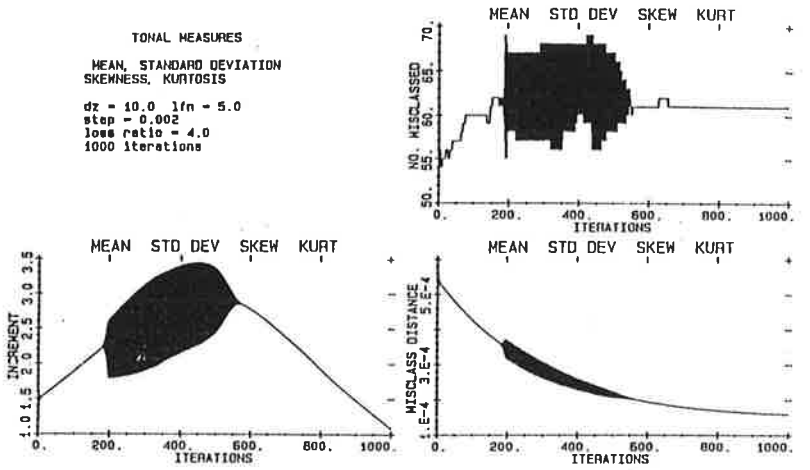

Figure 9.5g Loss ratio = 4.0.

The results in table 9.3 show a ratio of false negatives to false positives of 3:1 for a loss ratio of 1.0 this changes to a ratio of 1:3 for a loss ratio of 4.0. The optimum loss ratio is difficult to calculate at this stage because its precise effects on the later stages of processing are impossible to predict with accuracy. It is dependent upon the effectiveness of the feature extraction algorithms to identify false positives as clear areas and to extend the boundaries of feature areas to encompass the whole feature (which may include parts of false negative areas).

9.1.4 VARIATION OF THE ITERATION STEP SIZE

The step is the parameter which determines the the amount the decision boundry is changed with each iteration. Figures 9.6(a-e) show the results with a variation in step size from step = 0.001 to 0.005 with all other parameters held constant. The most striking feature is in the plots of the number of samples misclassified which show identical curves except that as the step is increased the curve falls into a proportionally smaller number of iterations. For example, the curve for step = 0.001 for 1000 iterations (figure 9.6a) is repeated for step = 0.002 but in 500 iterations (figure 9.6b). Similarly between figures 9.6b and 9.6d.

As the step is increased the maximum increment increases until it reaches 2.7 in figure 9.6d with a step of 0.004. Note the small amount of oscillation in the increment curve at the beginning of the iteration. When the step is increased to 0.005 in figure 9.6e more of the oscillating behaviour associated with a large increment value is seen.

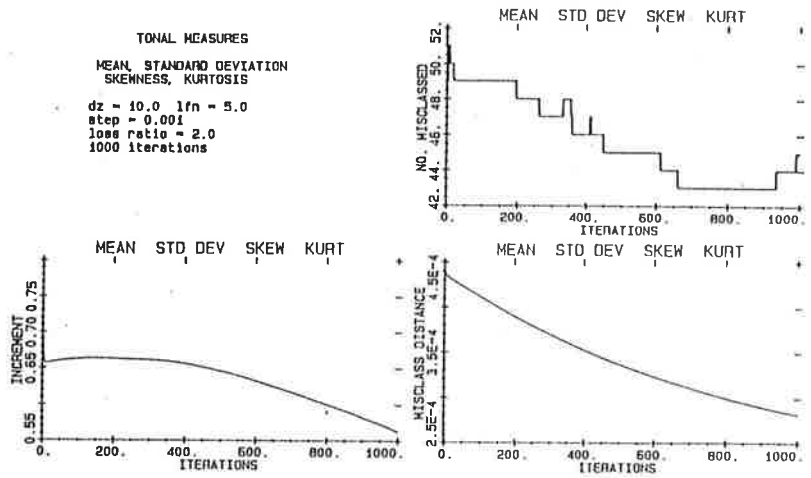

Figure 9.6a Step = 0.001.

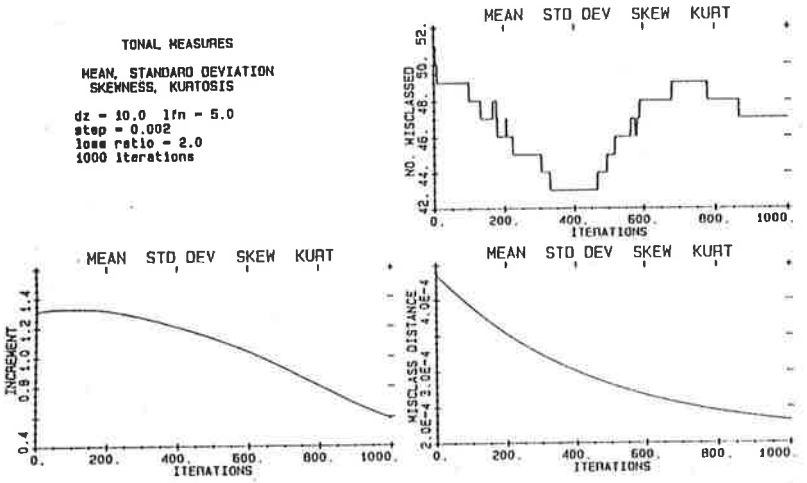

Figure 9.6b Step = 0.002.

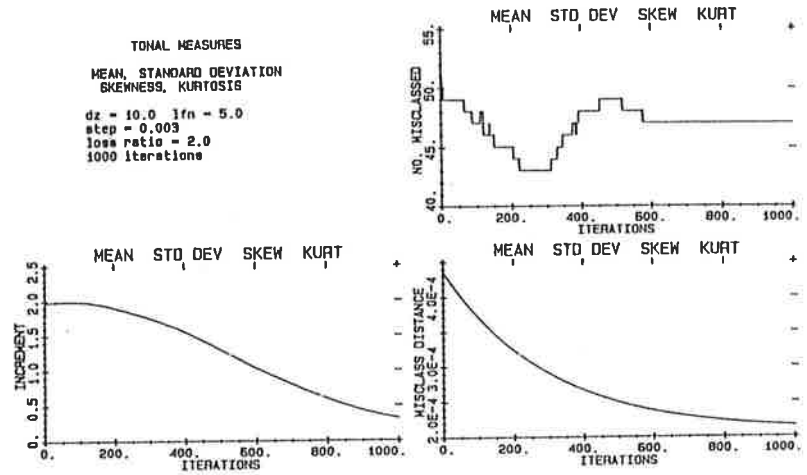

Figure 9.6c Step = 0.003.

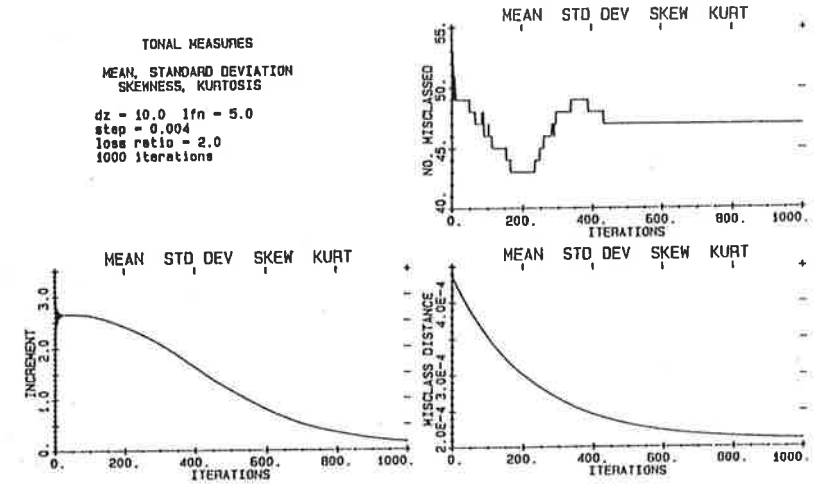

Figure 9.6d Step = 0.004.

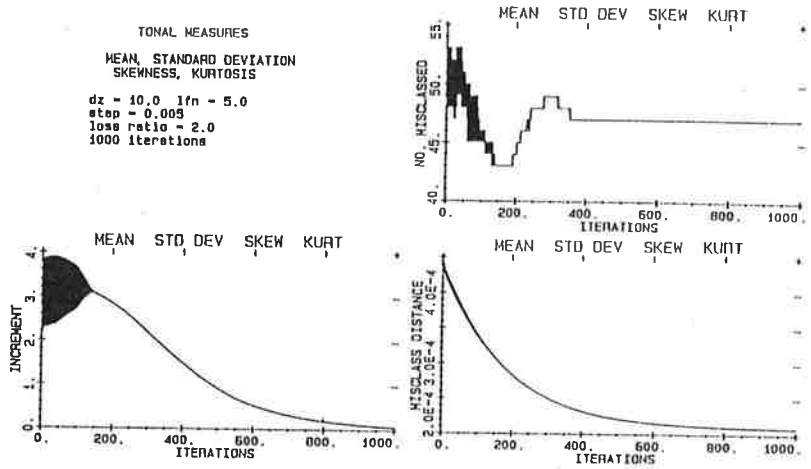

Figure 9.6e Step = 0.005.

The results of the variation of the step indicate that reduction of the step size can be used to eliminate oscillatory behaviour. The same final decision boundary can be achieved but only after many more iterations. The relationship between step size and the number of iterations is inversely linear, that is, if the step is halved the algorithm will need to do twice the iterations to reach the same final weight vector (as long as the iteration remains stable).

9.1.5 CONCLUSIONS OF ALGORITHM PARAMETER VALUES

A suitable choice of parameter values is necessary to prevent unstable and oscillatory behaviour of the iteration. The loss function and dead zone parameters must be balanced so that sufficient samples are included in the iteration and that the influence of the samples involved is neither too great nor too small. Values of lfn = 5.0 and dz = 10.0 are chosen as values that provide a stable iteration.

The loss ratio provides a measure of control by allowing the **a priori** perceived cost difference between false negatives and false positives to be considered in the iteration. Based upon an understanding that false negatives result in features passing undetected while false positives represent a computational burden a value of loss function = 2.0 is chosen.

The step determines the speed at which the iteration proceeds and so the time taken to reach a solution. A value of step = 0.002 is chosen as a value that provides a fast iteration while keeping the increment below 1.5, above which instability is likely.

9.2 USING COMBINATIONS OF TONAL MEASURES IN THE CLASSIFICATION ALGORITHM

Figures 9.7(a-o) show the algorithm applied to all the combinations of the four tonal measures in an attempt to determine the usefulness of each measure. All the plots are taken with the same initial position of coefficient values and parameter values except figures 9.7(l-o) which proved unstable with a step of 0.002 and so used a value of step = 0.001 and 4000 iterations. All the plots are well behaved with peaked and decreasing increments and flattening misclassification distance curves. The results are tabulated in table 9.4.

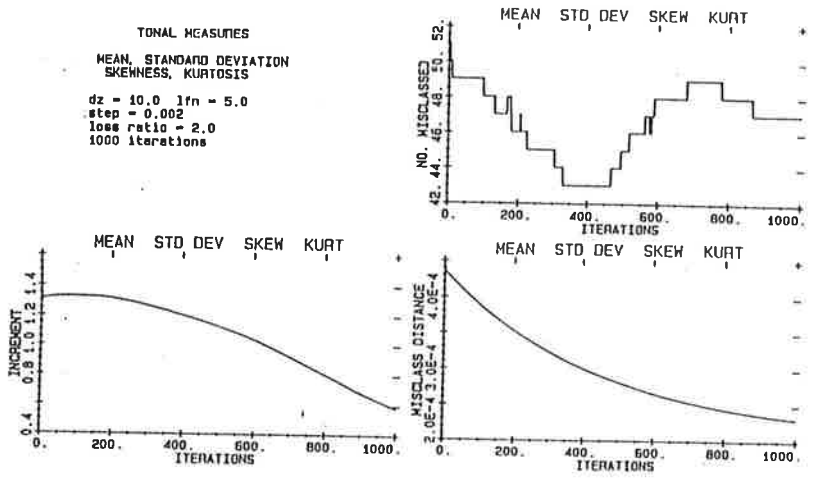

Figure 9.7a Mean, Std. Dev., Skewness, Kurtosis.

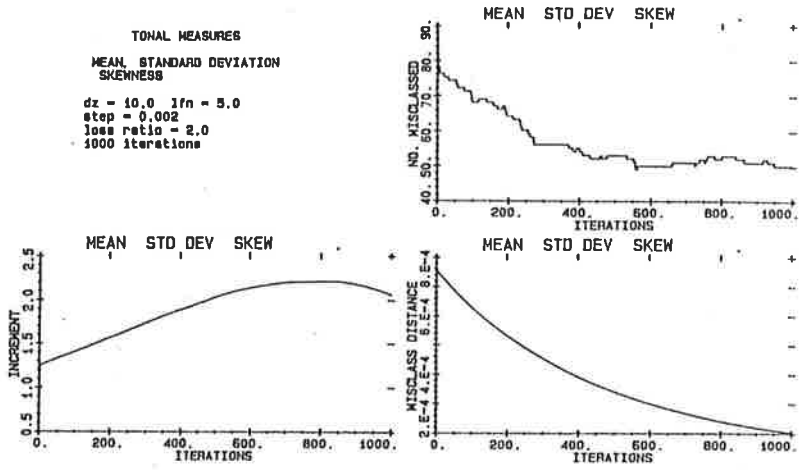

Figure 9.7b Mean, Std. Dev., Skewness.

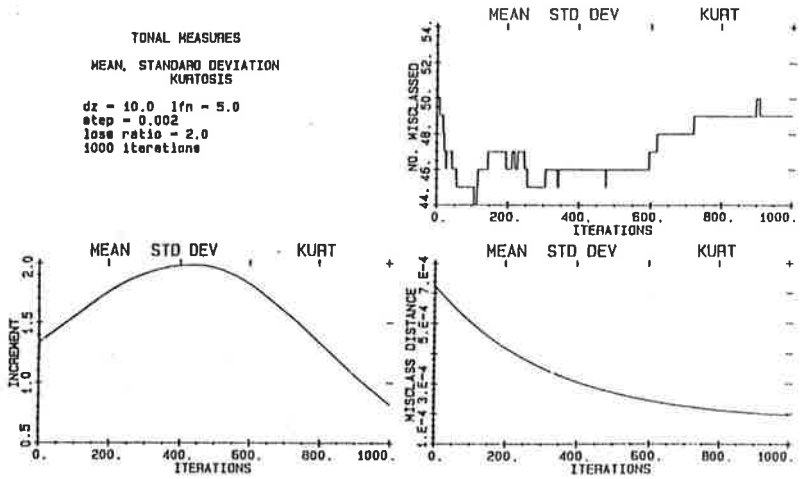

Figure 9.7c Mean, Std. Dev., Kurtosis.

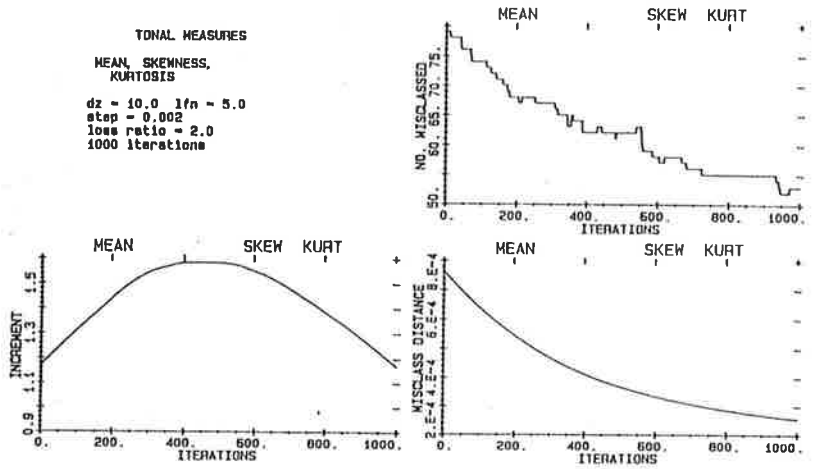

Figure 9.7d Mean, Skewness, Kurtosis.

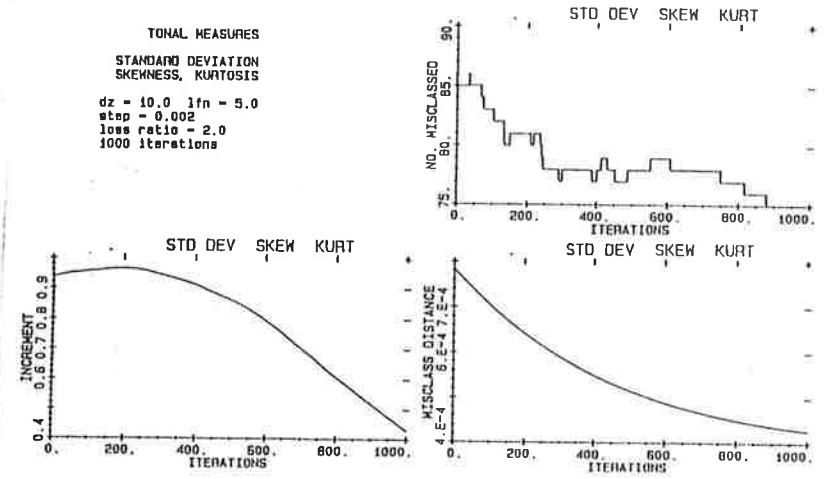

Figure 9.7e Std. Dev., Skewness, Kurtosis.

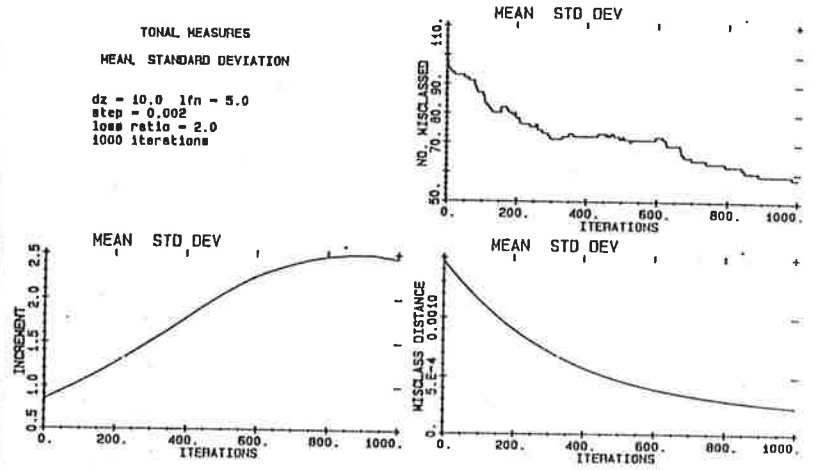

Figure 9.7f Mean, Std. Dev.

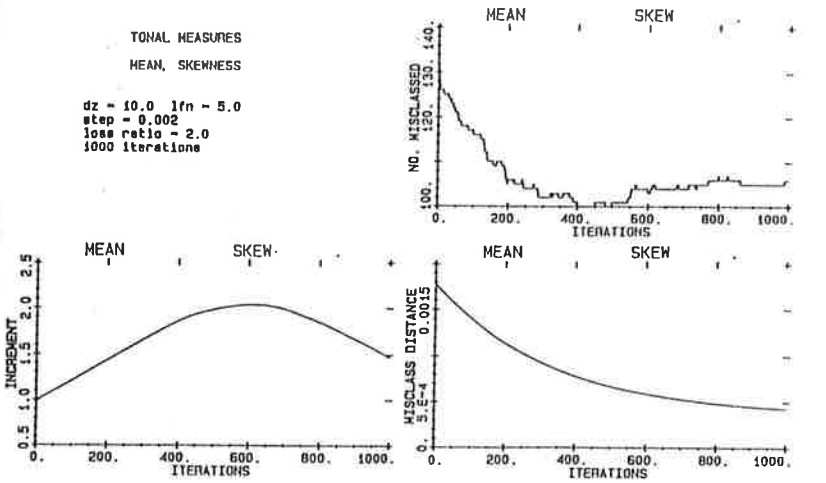

Figure 9.7g Mean, Skewness.

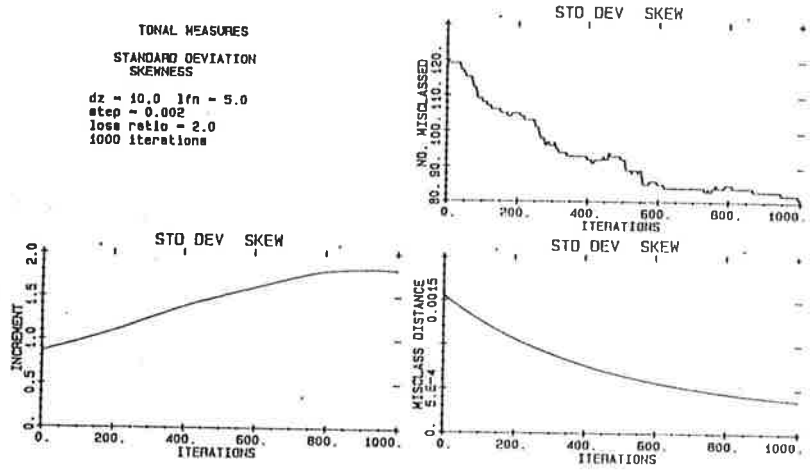

Figure 9.7h Std. Dev., Skewness.

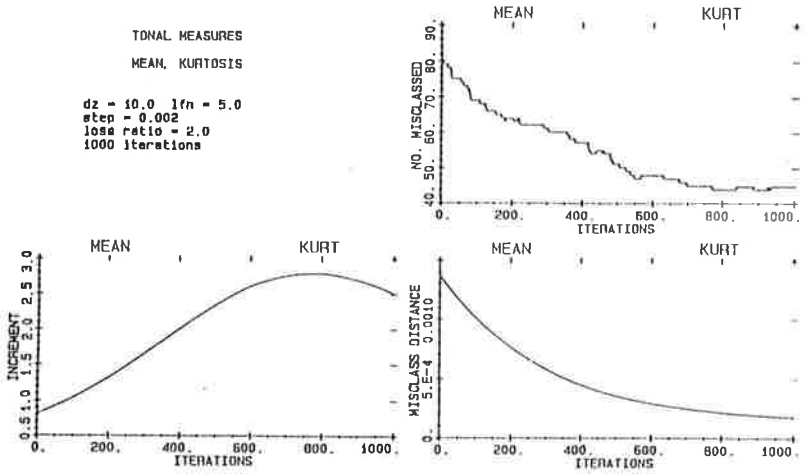

Figure 9.7i Mean, Kurtosis.

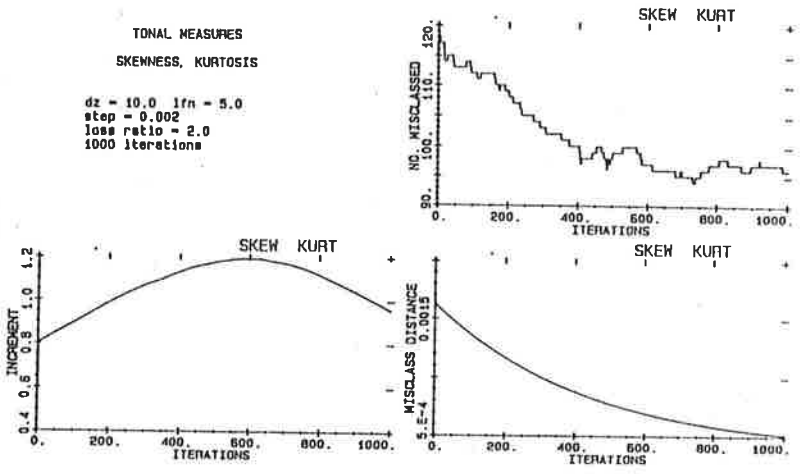

Figure 9.7j Skewness, Kurtosis.

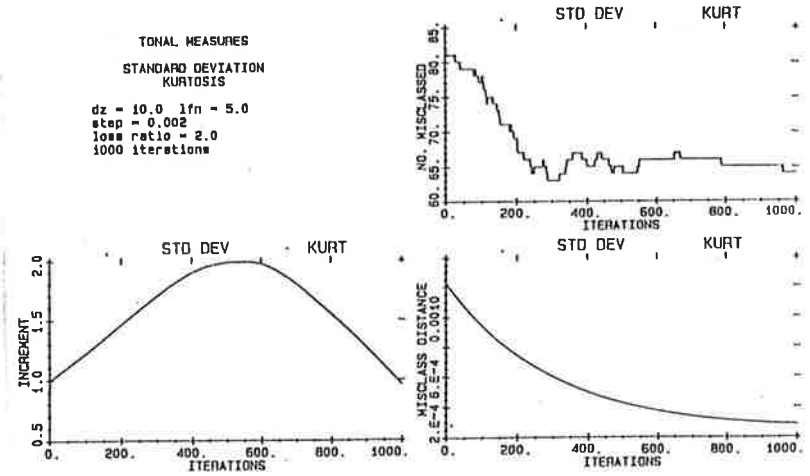

Figure 9.7k Std. Dev., Kurtosis.

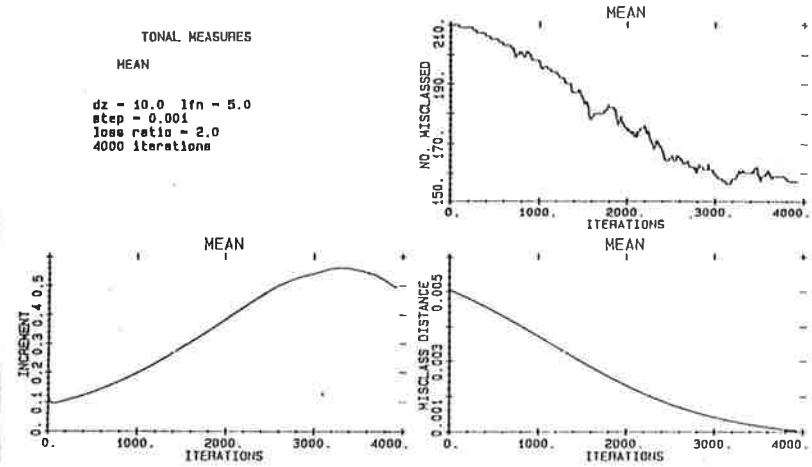

Figure 9.71 Mean.

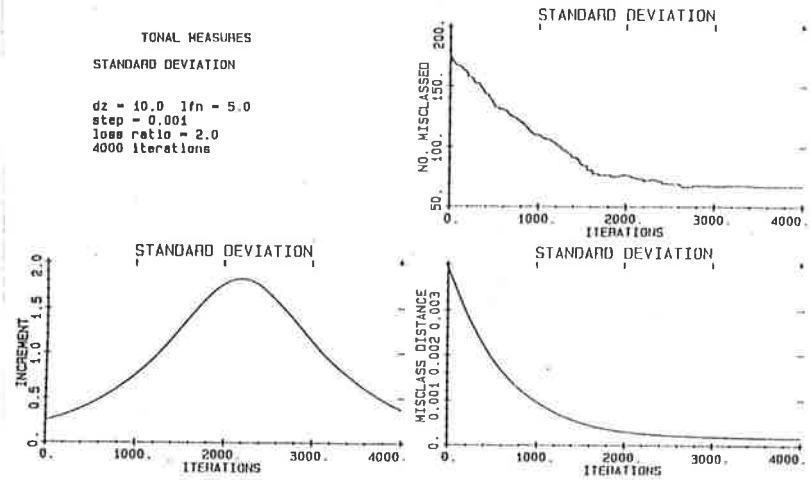

Figure 9.7m Std. Dev.

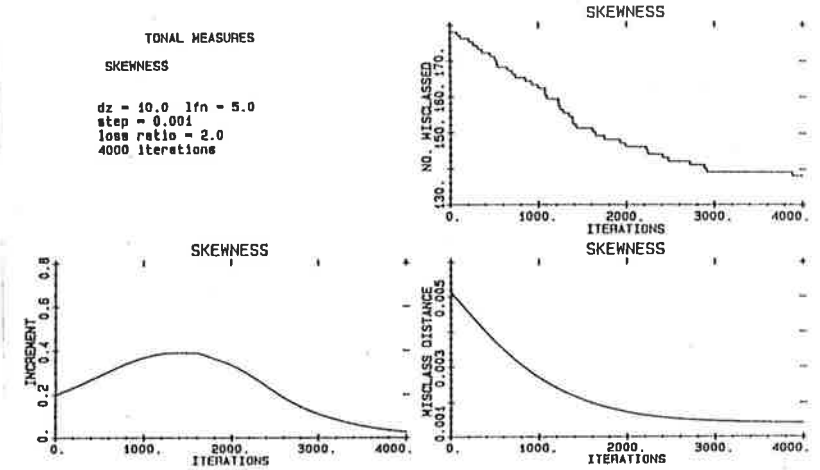

Figure 9.7n Skewness.

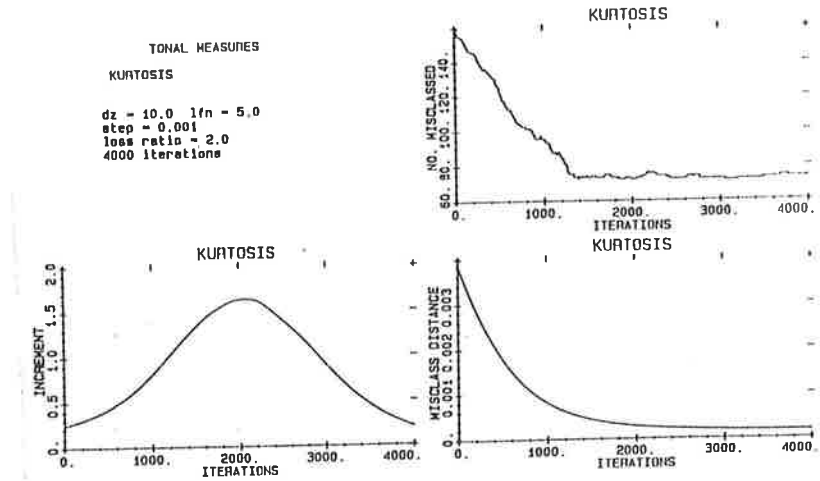

Figure 9.7o Kurtosis.

FEATURE MEASURES	MISCLASS DISTANCE ($*10^{-6}$)	NUMBER SAMPLES MISCLASSED	FALSE NEG	FALSE POS
(a) Mean, Std Dev, Skewness, Kurtosis	231.3	47	29	18
(b) Mean, Std Dev, Skewness	202.9	50	35	15
(c) Mean, Std Dev, Kurtosis	195.6	49	29	20
(d) Mean, Skewness, Kurtosis	254.3	53	43	10
(e) Std Dev, Skewness, Kurtosis	426.4	75	51	24
(f) Mean, Std Dev	236.3	58	43	15
(g) Mean, Skewness	423.8	106	82	24
(h) Std Dev, Skewness	365.8	81	64	17
(i) Mean, Kurtosis	174.6	45	34	11
(j) Skewness, Kurtosis	512.2	96	74	22
(k) Std Dev, Kurtosis	283.2	64	41	23
(l) Mean*	1016.2	157	97	60
(m) Std Dev*	161.9	67	43	24
(n) Skewness*	1411.6	138	138	0
(o) Kurtosis*	167.2	73	43	30

Table 9.4 The performance of combinations of feature measures to separate the two classes of the training set. dz = 10.0; lfn = 5.0; loss ratio = 2.0; step = 0.002; 1000 iterations; (* 4000 iterations). 1073 samples; 239 feature areas; 834 clear areas.

The decision boundaries that result after the iteration procedure are listed below:

w = (127.967, -140.576, 101.227, -139.050, 5.544)	9.1a
w = (145.662, -209.257, 136.371, 0.0, 5.417)	9.1b
w = (131.359, -161.700, 0.0, -175.397, 6.015)	9.1c
w = (147.423, 0.0, 111.762, -187.356, 5.806)	9.1d
w = (0.0, -142.218, 100.900, -138.325, 4.648)	9.1e
w = (142.896, -247.832, 0.0, 0.0, 5.596)	9.1f
w = (190.702, 0.0, 182.724, 0.0, 6.163)	9.1g
w = (0.0, -215.237, 129.703, 0.0, 4.621)	9.1h
w = (156.074, 0.0, 0.0, -256.309, 6.247)	9.1i
w = (0.0, 0.0, 116.286, -192.238, 4.658)	9.1j
w = (0.0, -179.546, 0.0, -186.625, 5.233)	9.1k
w = (214.522, 0.0, 0.0, 0.0, 5.596)	9.1l
w = (0.0, -365.634, 0.0, 0.0, 5.038)	9.1m
w = (0.0, 0.0, 183.370, 0.0, 6.002)	9.1n
w = (0.0, 0.0, 0.0, -349.833, 5.137)	9.1o

Recall that the samples in pattern space are of the form:

$\mathbf{x}$ = (mean, std dev, skewness, kurtosis, aug),

and that the decision boundary is the locus of all points such that:

$$\mathbf{w} \bullet \mathbf{x} = 0.$$

The results of using a single measure to derive a zero dimensional decision boundary in one dimensional pattern space are found in table 9.4(l,m,n,o). This is otherwise known as a threshold, with all the samples with a value of the measure greater than the threshold value being placed in one class and all those with a value less than the threshold being placed in the other class.

As mentioned previously care must be exercised when making comparisons of the misclassification distance between different iterations. The value of R varies with the dead zone

(constant in all these plots), the loss function (also constant), and the number of feature measures used (variable). Thus, (b), (c), (d) and (e) may be compared; (f), (g), (h), (i), (j) and (k) may be compared; and (l), (m), (n), and (o) may be compared.

If the misclassification distance, R, is used as a criterion on which to choose optimum combinations of feature measures then the single measures of standard deviation and kurtosis (m, o) are superior to those of mean and skewness (l,n). Examining combinations of two measures the most outstanding is mean and kurtosis (i), followed by mean and standard deviation (f) and then standard deviation and kurtosis (k). The three combinations that involve skewness (g, h, j) are all inferior. Examining combinations of three measures the best result is (c) which excludes skewness. The poorest result is (e) which excludes the mean.

Using the number of samples misclassified as a criterion yields identical results and also allows comparison across the whole table. The lowest number of misclassified samples is achieved using mean and kurtosis (i, 45). Similarly low values are achieved with (a, 47), (b, 50), and (c, 49).

Examining table 9.4 one can see that the measures of standard deviation and kurtosis, either alone or in combination, result in a low number of misclassifications except when in combination with skewness. Standard deviation and kurtosis together (k) gives a result of 64 misclassifications which is only marginally better that with either one alone (m, 67; o, 73). This suggests that the standard deviation and kurtosis are to an extent correlated. This can be visualised by examining the samples in the four dimensional pattern space (figure A2.5a and A2.5b in Appendix 2). The samples lie close to a line that bisects the space through the origin although the scatter gets larger at higher values. A sample that has a high value of standard deviation will also have a high value of kurtosis. Feature areas predominantly have high values of both measures.

The mean measure is an essential measure. By itself a mean threshold is not very effective as a discriminator (1, 157)

but its addition to other combinations of measures improves performance considerably. Adding mean to (k) to form (c) reduces the number of misclassified samples from 64 to 49. Similarly from (j) to (d) reduces the number of misclassified from 96 to 53, and from (h) to (b) the reduction is from 81 to 50. Clearly the mean adds important information concerning the presence of a feature that is not evident unless it is combined with other measures. The reason for the poor performance of the mean as a single threshold discriminator is attributed to the wide range of overall brightness levels of the wood itself and the variable gain inherent in the camera. Examination of the spread of samples in the mean dimension reveals that most clear areas have a large value of mean and feature areas have low values. There is a large amount of overlap, however, that can only be clarified with the presence of another measure.

Using the criteria of the misclassification distance, R, and the number of misclassified samples, M, leads to a choice of four combinations of tonal measures that are considered to give good results. These are (i), (a), (c), and (b) in decreasing order of effectiveness.

Using threshold values with mean or skewness results in a large number of samples misclassified. A large proportion of these misclassified samples are false negatives, that is, a feature area that is wrongly classified as clear. The skewness threshold results in all the misclassified areas being false negatives. Looked at another way the skewness threshold results in no false positives and so the few samples that are classified as feature areas are all true positives. The number of samples that it classifies as feature areas in this way is 239 - 138 = 101, or only 42.3% of the total number of 239 feature areas in the sample set. This presentation of the results is more enlightening that that of table 9.4 and is presented for all the tonal measure combinations using the training set of samples in table 9.5. The decision boundaries derived for these combinations using the training set are applied to the test set and presented in table 9.6.

TRAINING SET OF SAMPLES

FEATURE MEASURE COMBIN-ATION	NUMBER OF SAMPLES CLASSED AS FEATURE AREAS	NUMBER OF THE 239 FEATURE AREAS CORRECTLY CLASSIFIED No.	%	PERCENTAGE OF AREAS CLASSED AS FEATURES THAT ARE TRUE POSITIVE
(a)	228	210	87.9	92.1
(b)	219	204	85.4	93.2
(c)	230	210	87.9	91.3
(d)	206	196	82.0	95.1
(e)	212	188	78.7	88.7
(f)	211	196	82.0	92.9
(g)	181	157	65.7	86.7
(h)	192	175	73.2	91.1
(i)	216	205	85.8	90.5
(j)	187	165	69.0	88.2
(k)	221	198	82.8	89.6
(l)	202	142	59.4	70.3
(m)	220	196	82.0	89.1
(n)	101	101	42.3	100.0
(o)	226	196	82.0	86.7

Table 9.5 Combinations of tonal measures are used to derive decision boundaries with the training set. These figures can be used to describe the confidence of the decision boundary to identify the 239 feature areas.

TEST SET OF SAMPLES

FEATURE MEASURE COMBIN-ATION	NUMBER OF SAMPLES CLASSED AS FEATURE AREAS	NUMBER OF THE 189 FEATURE AREAS CORRECTLY CLASSIFIED No.	%	PERCENTAGE OF AREAS CLASSED AS FEATURES THAT ARE TRUE POSITIVE
(a)	197	176	93.1	89.3
(b)	185	169	89.4	91.4
(c)	195	174	92.1	89.2
(d)	196	172	91.0	87.8
(e)	185	161	85.2	87.0
(f)	164	151	79.9	92.1
(g)	173	159	84.1	91.9
(h)	161	150	79.4	93.2
(i)	197	175	92.6	88.8
(j)	173	145	76.7	83.8
(k)	188	164	86.8	87.2
(l)	136	104	55.0	76.5
(m)	157	143	75.7	91.1
(n)	102	96	50.8	94.1
(o)	199	163	86.2	81.9

Table 9.6 Decision boundaries derived from the training set applied to the test set. There are a total of 189 feature areas.

These figures present clearly the factors that are important for the detection of features in wood. It is imperative that the maximum number of feature areas are correctly classed as features. The central purpose of this study is to determine the feature measures that will best enable the detection of feature areas. A corresponding aspect is the confidence that the feature areas detected are true positives. False positives have to be

detected in later stages of processing and so it is desirable to reduce their presence as early as possible.

The feature measure combinations that result in the detection of the largest proportion of the 239 feature areas of the training set are, in order of effectiveness, (a), (c), (i) and (b). These are the same four combinations that were chosen as optimal from table 9.4 based upon the misclassification distance and the number of samples misclassified. The areas classed as feature areas contain more than 90% of true positives.

When the derived decision boundaries are applied to the test set of samples the results are in strong agreement. The same four combinations perform well with the exeption of (b) which is displaced by (d). A slightly greater proportion of the feature areas are classed correctly (for example, in (a) 93.1% of the feature areas are correctly classified in the test set as opposed to 87.9% in the training set) while the proportion of the areas classed as feature areas that are true positives is generally slightly less. This agreement between the results derived from the training set and applied to the test set increases the confidence that the derived decision boundaries are not artifacts of the training set.

These results lead to the following conclusions:

a) The tonal measure of mean is vital to the identification of feature areas in radiata pine.

b) The measures of standard deviation and kurtosis are correlated and describe similar features.

c) Best results are achieved when they are both used together with the mean.

d) The measure of skewness is of marginal benefit and its inclusion or exclusion has little effect.

9.2.1 EXAMINATION OF LOCAL AREAS MISCLASSIFIED BY TONAL MEASURES

An examination of the misclassified samples is performed to understand the types of areas that fail to be classed correctly by decision boundaries derived with different tonal measure combinations. Using the decision boundary derived using

the four tonal measures, 9.1a, the distance of each misclassified area is listed and a description given. The reader can locate the local areas by referring to the WOOD library in Appendix 1A.

```
w       = (128.0, -140.6, 101.2, -139.1, 5.54)
mean    = (185.21,    15.32,  -18.08,    27.16)
std dev = (823.85,   341.24, 3053.65,   638.85)
```

IMAGE	LOCAL AREA	w.x	CLASSIFICATION
wood01	3 3	9.995	FALSE NEG
	Edge of large intergrown knot		
wood01	3 4	1.529	FALSE NEG
	Centre of large intergrown knot		
wood01	3 5	8.814	FALSE NEG
	Edge of large intergrown knot		

All three areas are of the same knot and are connected to correctly classified area of bark

wood04	2 3	-0.341	FALSE POS
	Very small portion of encased knot		
wood05	0 5	-2.778	FALSE POS
	Dark growth rings		
wood06	2 4	-1.838	FALSE POS
	Growth rings near large spike knot		
wood06	5 5	-3.218	FALSE POS
	Growth rings near large spike knot		
wood08	3 4	1.974	FALSE NEG
	Edge of intergrown knot - connected		
wood09	4 1	-2.898	FALSE POS
wood09	5 1	-0.152	FALSE POS
wood09	6 1	-5.075	FALSE POS

All three areas contain a portion of resin - could be considered to be true positives

wood09	6 2	-0.546	FALSE POS
	Dark growth rings and resin		
wood10	2 4	-4.067	FALSE POS
	Very small portion of encased knot		

wood10 4 5 -0.678 FALSE POS

wood10 5 2 -0.285 FALSE POS

wood10 5 5 -3.421 FALSE POS

wood10 6 3 -0.491 FALSE POS

All four areas have dark growth rings and specular reflections

wood11 1 6 2.086 FALSE NEG

Edge of resin pocket - connected areas of resin correctly classed

wood12 2 2 -1.093 FALSE POS

Resin stain - connected area with darker stain is correctly classed as feature area

wood12 4 1 6.256 FALSE NEG

wood12 5 0 6.535 FALSE NEG

wood12 6 1 1.372 FALSE NEG

All three are portions of a large, faint intergrown knot - only part of this knot is correctly classified

wood12 4 5 4.351 FALSE NEG

wood12 4 6 2.768 FALSE NEG

Both contain the edge of a partially encased knot - the bark encasement is correctly classified - connected

wood13 2 2 10.010 FALSE NEG

wood13 2 3 1.935 FALSE NEG

wood13 3 2 3.093 FALSE NEG

wood13 3 3 0.884 FALSE NEG

The four areas encompass the whole area of an intergrown knot

wood14 2 5 0.878 FALSE NEG

Edge of small knot - connected

wood14 3 2 5.724 FALSE NEG

wood14 4 0 3.411 FALSE NEG

wood14 4 2 1.163 FALSE NEG

All three areas are part of a very faint intergrown knot of which a connected area of bark is classified correctly

wood15 1 3 -2.538 FALSE POS

Dark growth rings

wood15 2 3 -0.191 FALSE POS

Growth ring near knot

wood15 2 1 3.528 FALSE NEG

wood15 3 0 5.331 FALSE NEG

Both areas are part of a faint intergrown knot

wood17 1 5 3.702 FALSE NEG

Edge of intergrown knot - connected

wood17 2 2 2.521 FALSE NEG

Edge of faint intergrown knot - connected

wood17 3 3 1.947 FALSE NEG

Centre of faint intergrown knot - connected

wood19 2 1 2.478 FALSE NEG

Edge of partially encased knot - connected

wood21 1 3 3.872 FALSE NEG

Growth rings of diffuse edge of intergrown spike knot

wood21 2 2 1.555 FALSE NEG

Small portion of dark growth rings associated with spike knot

wood22 2 3 -0.825 FALSE POS

Growth ring distortion due to proximity to spike knot

wood24 2 2 0.724 FALSE NEG

Centre of faint intergrown knot - connected

wood26 3 0 4.157 FALSE NEG

Edge of faint intergrown knot - connected

wood38 0 1 -0.148 FALSE POS

Very small portion of intergrown knot

wood38 3 6 1.141 FALSE NEG

Growth rings of large intergrown knot - connected - could be classed as a true negative

No OF MISCLASSED = 47

No OF FALSE NEGATIVES = 29

No OF FALSE POSITIVES = 18

Problem areas are the interior and edges of large knots that are not very dark. There is not sufficient density to excite the mean nor sufficient change to excite the variance or kurtosis. A similar explanation applies to the intergrown pin

knots and the edges of diffuse spike knots. The knot in the upper right portion of WOOD12 is an example of a very faint knot that is impossible to detect with tonal measures. The large values for the misclassification distance testify to this.

A number of large knots have the central part of the knot misclassified but they often contain bark on one side that is correctly classed. The presence of the knot can be inferred by the presence of this bark which is usually longer across the width of the board than along the length of the board. This bark is due to the encasement of material trapped in the angle of steep branches. Examples are WOOD01, WOOD12, and WOOD14.

A number of areas could be given alternate classifications. WOOD09 contains three areas with resin that are labelled positive by the decision boundary and WOOD12 also contains an area with a portion of resin. WOOD38 contains growth rings associated with a large intergrown spike knot and is labelled by the decision boundary as a negative. It is connected to a number of true positives and does not influence the definition of the feature.

A large number of the false negatives are part of a feature of which the rest is correctly identified. Sometimes this results in only half of a small feature being defined. If the false negative is part of a large feature it is usually the edge and its absence is not of any consequence if the extent of the feature is defined by other correctly classed areas.

An extensive analysis of the misclassified local areas is of limited use at this stage. When the feature local areas are joined to form feature areas then the influence of the misclassified areas will become more apparent. A more beneficial study of their effect on the discrimination of features is found in chapter 10.

9.3 THE ADDITION OF TEXTURE MEASURES TO IMPROVE THE CLASSIFICATION OF LOCAL AREAS

Given the proven capabilities of the tonal measures to identify feature areas there only remains the task of improving

this to overcome the errors made by the tonal dichotomy. These have been identified as the faint intergrown knots that do not present sufficient contrast to be detected by tonal measures. A characteristic of intergrown knots is that they disrupt the appearance of the growth rings to a greater extent than an encased knot. This suggests that texture measures may be useful in identifying those faint features that are misclassified as clear with the use of tonal measures.

To examine this idea the texture measures of the WOOD library of images are calculated. The measures of contrast (CON), angular second moment (ASM), correlation (COR), entropy (ENT), local homogeniety (LHOM), cluster shade (CSH), and cluster prominence (CPR) are calculated for each local area. Unlike the previous work on the FEAT set of images which used an unequal probability quantising technique the reduction of grey levels from 256 to 16 is performed using an equal probability quantising method for the WOOD set of images. This is done to be consistent with previous published studies of texture. The value of each measure is the average of the measure in the four directions, for example:

$$CON_{tdv} = 1/4(CON(P_0(tdv)) + CON(P_{45}(tdv)) + CON(P_{90}(tdv)) + CON(P_{135}(tdv)))$$

Eight values of the texture displacement vector (tdv) are examined: tdv = 1, 2, 3, 4, 6, 8, 12, 16.

Each texture measure is added singly to the four tonal measures and the classification iteration performed. This is performed several times with different initial texture measure coefficients in the weight vector (this texture measure coefficient is referred to as tex). The three possible outcomes are:

a) The coefficient is unsuitable and the iteration will not reach a position that is better than that using just the tonal measures.

b) The initial coefficient is reasonably close to the optimum value but the measure contains no new information. The final value of the coefficient will tend to zero and so it will have no

effect on the ability of the tonal measures to classify the samples.

c) If the texture measure does contain additional information it will result in a lower number of samples misclassified and a lower misclassification distance. This is demonstrated in table 9.7 with the contrast texture measure in the 0° direction.

INITIAL w = (100.0,-100.0, 100.0,-100.0,-100.0, 5.00) FINAL w = (127.4,-138.6, 104.7,-134.9, -87.7, 5.84) R = 0.000202 M = 50 false pos = 27 false neg = 23
INITIAL w = (100.0,-100.0, 100.0,-100.0, -75.0, 5.00) FINAL w = (131.2,-144.2, 103.6,-140.0, -72.8, 5.76) R = 0.000192 M = 47 false pos = 25 false neg = 22
INITIAL w = (100.0,-100.0, 100.0,-100.0, -50.0, 5.00) FINAL w = (132.6,-146.8, 102.1,-142.5, -61.5, 5.70) R = 0.000191 M = 45 false pos = 23 false neg = 22
INITIAL w = (100.0,-100.0, 100.0,-100.0, -25.0, 5.00) FINAL w = (132.9,-147.9, 101.1,-143.9, -51.0, 5.64) R = 0.000192 M = 43 false pos = 20 false neg = 23
INITIAL w = (100.0,-100.0, 100.0,-100.0, 25.0, 5.00) FINAL w = (131.6,-145.5, 101.2,-143.6, -25.0, 5.56) R = 0.000205 M = 45 false pos = 18 false neg = 27
INITIAL w = (100.0,-100.0, 100.0,-100.0, 50.0, 5.00) FINAL w = (128.0,-139.2, 102.5,-139.0, -3.5, 5.53) R = 0.000230 M = 47 false pos = 18 false neg = 29
INITIAL w = (100.0,-100.0, 100.0,-100.0, 75.0, 5.00) FINAL w = (115.7,-121.1, 103.8,-122.2, 41.1, 5.51) R = 0.000332 M = 54 false pos = 13 false neg = 41
INITIAL w = (100.0,-100.0, 100.0,-100.0, 100.0, 5.00) FINAL w = (100.1,-100.2, 100.1,-100.2, 99.7, 5.82) R = 0.000532 M = 84 false pos = 15 false neg = 69

Table 9.7 Deriving the optimum initial coefficient for the texture measure contrast in the 0° direction. tdv = 12; lfn = 5.0; dz = 10.0; step = 0.002; loss ratio = 2.0; 2000 iterations.

The initial coefficient for the contrast texture measure is chosen to be tex = -50 because, as well as giving a low value

for both R and M, the final value of the coefficient is very close. If an initial value of tex = -75 is used the final value is less at tex = -72.8. If an initial value of tex = -25 is used the final value is tex = -51.0 indicating that the final value that is approached in the iteration is between -50 and -75.

For initial values in the range from tex = 25 to tex = -100 the final values of the decision boundary coefficients are very similar to that derived for the tonal measures alone (equation 9.1a). As the initial coefficient of the contrast measure becomes more positive from tex = 50 to tex = 100 the final values of the other coefficients become less changed from the initial values. This indicates that the increment is too small to change the position of the decision boundary significantly. This can be rectified in several ways.

The first method is to choose a more suitable initial coefficient. This is the course adopted in this study. An alternative method is to increase the increment of the classification algorithm by increasing the step. Caution must be exercised to ensure that the iteration remains stable.

The procedure outlined for the calculation of the contrast texture measure coefficient is repeated for the other six texture measures. The initial decision boundary for each iteration is of the form:

w = (100.0, -100.0, 100.0, -100.0, tex, 5.0)

where the most suitable initial values of tex are estimated to be as follows:

CON - tex = -50.0
ASM - tex = 100.0
COR - tex = 75.0
ENT - tex = 100.0
LHOM - tex = 50.0
CSH - tex = -50.0
CPR - tex = 50.0

The iteration is performed 2000 times for the average of each texture measure over the four directions for each of the

eight values of the texture displacement vector, tdv, and the results are presented in table 9.8.

TONAL MEASURES No. MISCLASSIFIED, M = 48
MISCLASSIFICATION DISTANCE, R = 222 * 10^{-6}

tdv		1	2	3	4	6	8	12	16
TONAL + CON	M	47	47	47	47	47	47	47	49
	R	222	222	222	222	221	220	216	214
TONAL + ASM	M	76	75	66	61	56	54	54	52
	R	379	380	343	317	286	247	231	219
TONAL + COR	M	48	47	47	47	46	45	46	45
	R	257	257	258	257	249	241	231	229
TONAL + ENT	M	86	82	75	61	57	53	49	57
	R	475	423	375	329	291	258	242	230
TONAL + LHOM	M	47	47	47	47	47	47	46	46
	R	227	225	225	223	221	218	214	213
TONAL + CSH	M	47	47	47	46	45	45	42	40
	R	218	217	215	213	208	202	195	189
TONAL + CPR	M	47	47	47	47	45	45	42	40
	R	218	217	215	213	209	203	196	190

Table 9.8 The influence of the texture displacement vector, tdv, and the seven texture measures on the WOOD training set. lfn = 5.0; dz = 10.0; step = 0.002; loss ratio = 2.0; 2000 iterations.

The measures of ASM and ENT both have higher values of M and R than with the tonal measures alone except at high values of the tdv. Examination of the coefficients of the decision boundary generated in these iterations reveals that the change in the decision boundary at low values of tdv is very small, deviating little from the initial position. As the value of the tdv is

increased the change in the position of the final decision boundary increases until it reaches a position comparable to that achieved using only tonal measures.

The texture measures of CON, ASM, COR, ENT and LHOM do not achieve a final value of M or R that is a significant improvement over that achieved with tonal measures alone. Table 9.9 reveals what is happening with these iterations using the example of the correlation texture measure. For each value of tdv the texture coefficient decreases as the iteration proceeds while the tonal measure coefficients increase. Comparing the tonal measure coefficients with that obtained using only the tonal measures in equation 9.1a one can see that they approach similar values. The decision boundary iterates toward the same place in pattern space while the low value of the texture measure coefficient ensures a minimal influence of the texture measure. Similar results are achieved with contrast and local homogeniety except that the final values of the texture coefficient in these cases are of the order of 1 and 10 respectively. Angular second moment and entropy both exhibit decreasing values of the texture measure coefficient indicating that if a larger step were used or more iterations performed a similar result would follow.

INITIAL w(tdv) = (100.0,-100.0, 100.0,-100.0, 75.0, 5.00)	
FINAL w(1) = (123.2,-135.7, 103.6,-134.0, 25.6, 5.38) R = 0.000257 M = 48 false pos = 14 false neg = 34	
FINAL w(2) = (123.3,-135.7, 103.8,-134.1, 25.7, 5.39) R = 0.000257 M = 47 false pos = 14 false neg = 33	
FINAL w(3) = (123.3,-135.1, 104.1,-133.7, 25.8, 5.42) R = 0.000258 M = 47 false pos = 14 false neg = 33	
FINAL w(4) = (123.4,-134.9, 104.5,-133.7, 25.6, 5.44) R = 0.000257 M = 47 false pos = 13 false neg = 34	
FINAL w(6) = (124.5,-136.8, 104.9,-135.7, 24.3, 5.46) R = 0.000249 M = 46 false pos = 13 false neg = 33	
FINAL w(8) = (125.7,-138.9, 105.0,-137.7, 24.1, 5.48) R = 0.000241 M = 45 false pos = 13 false neg = 32	
FINAL w(12) = (127.1,-142.7, 104.6,-140.6, 27.9, 5.48) R = 0.000231 M = 46 false pos = 14 false neg = 32	
FINAL w(16) = (126.1,-143.6, 103.7,-140.7, 35.4, 5.44) R = 0.000229 M = 45 false pos = 14 false neg = 31	

Table 9.9 Decision boundaries derived for the correlation texture measure averaged over the four directions for different values of the texture displacement vector, tdv. lfn = 5.0; dz = 10.0; step = 0.002; loss ratio = 2.0; 2000 iterations.

The results of table 9.8 for the measures of cluster shade and cluster prominence show a significant improvement over the tonal measures at the higher values of tdv in the final values of both M and R. An examination of the values of the coefficients of the decision boundaries is instructive and is displayed in table 9.10 for the texture measure of cluster shade.

INITIAL w(tdv) = (100.0,-100.0, 100.0,-100.0, -50.0, 5.00)
FINAL w(1) = (129.5,-146.3, 101.1,-144.0, -2.3, 5.58) R = 0.000218 M = 47 false pos = 18 false neg = 29
FINAL w(2) = (129.5,-146.7, 101.3,-144.3, -2.9, 5.57) R = 0.000217 M = 47 false pos = 18 false neg = 29
FINAL w(3) = (129.7,-147.5, 101.9,-145.2, -5.0, 5.57) R = 0.000215 M = 47 false pos = 18 false neg = 29
FINAL w(4) = (130.0,-148.3, 102.4,-146.0, -7.1, 5.57) R = 0.000213 M = 46 false pos = 17 false neg = 29
FINAL w(6) = (130.5,-149.8, 102.9,-147.7, -11.3, 5.58) R = 0.000208 M = 45 false pos = 16 false neg = 29
FINAL w(8) = (131.1,-151.8, 103.3,-149.7, -17.7, 5.59) R = 0.000202 M = 45 false pos = 16 false neg = 29
FINAL w(12) = (131.7,-154.0, 103.3,-151.6, -27.4, 5.60) R = 0.000195 M = 42 false pos = 15 false neg = 27
FINAL w(16) = (131.8,-155.7, 103.0,-153.0, -36.4, 5.62) R = 0.000189 M = 40 false pos = 14 false neg = 26

Table 9.10 Decision boundaries derived for the **cluster shade** texture measure averaged over the four directions for different values of the texture displacement vector, tdv. lfn = 5.0; dz = 10.0; step = 0.002; loss ratio = 2.0; 2000 iterations.

Table 9.10 shows clearly that the texture measure coefficient tends to zero if the measure contains no new information to the classification but will tend toward a real value if the measure does contain information. Cluster shade and cluster prominence are concluded to both contain relevant information for values of tdv from 6 to 16.

9.3.1 INVESTIGATION OF THE ROLE OF TEXTURE DIRECTION

That some texture measures do not contain useful information is inconsistent with the perception of wood as a textured image until one examines the structure of the texture in more detail. The predominant texture is associated with the growth rings which appear as horizontal lines. These lines are

deflected by the presence of features but the separation between them is extemely variable, depending on the angle at which the plank is cut from the log. Attempting to detect grey level dependence in the 90° direction will be corrupted by this wide variation in growth ring spacing. However, there will be a strong grey level dependence in the 0° direction for clear wood where the growth rings are all horizontal. This dependence will be degraded when the growth rings are deflected by a feature.

In order to examine the direction dependence of the texture measures each of the four directions of each texture measure for each value of tdv is used in turn with the tonal measures. This is a total of 224 runs of 2000 iterations each which took the equivalent of 2 weeks continuous processing on an IBM-XT. The results are presented in tables 9.11 to 9.17.

CONTRAST

tdv		1	2	3	4	6	8	12	16
TONAL + CON$_{0°}$	M	40	41	40	40	41	40	45	49
	R	189	193	184	184	185	188	191	196
TONAL + CON$_{45°}$	M	48	48	47	47	47	47	47	47
	R	233	233	219	219	219	218	216	213
TONAL + CON$_{90°}$	M	49	49	47	47	47	47	47	46
	R	243	243	228	228	227	225	220	218
TONAL + CON$_{135°}$	M	49	49	47	47	47	47	47	47
	R	238	238	226	227	228	230	230	230

Table 9.11 The influence of the texture displacement vector, tdv, and the four directions of CONTRAST on the WOOD training set. lfn = 5.0; dz = 10.0; step = 0.002; loss ratio = 2.0; 2000 iterations; CON coefficient = -50.0.

ANGULAR SECOND MOMENT

tdv		1	2	3	4	6	8	12	16
TONAL + ASM$_{0}$ 0	M	46	49	46	47	49	47	49	48
	R	203	197	191	190	188	186	183	189
TONAL + ASM$_{45}$ 45	M	75	75	65	63	56	52	49	51
	R	487	487	393	344	302	266	243	243
TONAL + ASM$_{90}$ 90	M	99	97	97	92	90	80	67	60
	R	583	577	560	541	516	456	335	277
TONAL + ASM$_{135}$ 135	M	102	102	101	92	91	86	77	61
	R	554	554	547	535	520	496	423	303

Table 9.12 The influence of the texture displacement vector, tdv, and the four directions of **ANGULAR SECOND MOMENT** on the WOOD training set. lfn = 5.0; dz = 10.0; step = 0.002; loss ratio = 2.0; 2000 iterations; ASM coefficient = 100.0.

CORRELATION

tdv		1	2	3	4	6	8	12	16
TONAL + COR$_{0}$ 0	M	43	44	42	41	45	47	49	51
	R	193	197	188	188	189	191	191	196
TONAL + COR$_{45}$ 45	M	46	46	46	46	44	42	43	45
	R	265	265	245	245	242	234	226	222
TONAL + COR$_{90}$ 90	M	56	55	52	51	50	48	49	50
	R	317	321	289	286	272	261	243	242
TONAL + COR$_{135}$ 135	M	52	52	48	50	52	54	52	52
	R	295	295	279	287	295	297	294	288

Table 9.13 The influence of the texture displacement vector, tdv, and the four directions of **CORRELATION** on the WOOD training set. lfn = 5.0; dz = 10.0; step = 0.002; loss ratio = 2.0; 2000 iterations; COR coefficient = 75.0.

ENTROPY

tdv		1	2	3	4	6	8	12	16
TONAL + ENT$_{0°}$	M	43	46	45	46	48	49	44	46
	R	205	201	194	192	189	186	186	193
TONAL + ENT$_{45°}$	M	78	78	67	63	59	54	47	51
	R	481	481	394	348	302	265	240	240
TONAL + ENT$_{90°}$	M	98	102	98	94	90	81	64	59
	R	576	569	553	536	504	409	315	270
TONAL + ENT$_{135°}$	M	98	98	95	89	92	90	82	66
	R	545	545	540	529	515	491	422	325

Table 9.14 The influence of the texture displacement vector, tdv, and the four directions of **ENTROPY** on the WOOD training set. lfn = 5.0; dz = 10.0; step = 0.002; loss ratio = 2.0; 2000 iterations; ENT coefficient = 100.0.

LOCAL HOMOGENIETY

tdv		1	2	3	4	6	8	12	16
TONAL + LHOM$_{0°}$	M	41	41	41	42	42	45	51	52
	R	184	185	179	179	182	187	195	198
TONAL + LHOM$_{45°}$	M	49	49	49	48	47	47	47	45
	R	239	239	224	223	221	219	214	212
TONAL + LHOM$_{90°}$	M	47	47	46	46	46	46	47	48
	R	260	262	243	241	238	233	226	223
TONAL + LHOM$_{135°}$	M	48	48	46	46	46	46	48	49
	R	258	258	244	245	245	245	243	241

Table 9.15 The influence of the texture displacement vector, tdv, and the four directions of **LOCAL HOMOGENIETY** on the WOOD training set. lfn = 5.0; dz = 10.0; step = 0.002; loss ratio = 2.0; 2000 iterations; LHOM coefficient = 50.0.

CLUSTER SHADE

tdv		1	2	3	4	6	8	12	16
TONAL + CSH$_0$ 0	M	49	46	43	39	38	38	41	42
	R	229	219	195	187	181	182	188	193
TONAL + CSH$_0$ 45	M	49	49	47	45	45	45	41	40
	R	230	230	213	210	207	201	196	192
TONAL + CSH$_0$ 90	M	49	49	47	47	47	47	46	44
	R	237	239	224	224	221	218	213	211
TONAL + CSH$_0$ 135	M	49	49	47	48	47	46	47	42
	R	231	231	217	216	215	211	206	199

Table 9.16 The influence of the texture displacement vector, tdv, and the four directions of **CLUSTER SHADE** on the WOOD training set. lfn = 5.0; dz = 10.0; step = 0.002; loss ratio = 2.0; 2000 iterations; CSH coefficient = -50.0.

CLUSTER PROMINENCE

tdv		1	2	3	4	6	8	12	16
TONAL + CPR$_0$ 0	M	49	45	40	38	38	39	42	45
	R	230	218	194	186	181	183	188	193
TONAL + CPR$_0$ 45	M	49	49	47	45	45	44	42	41
	R	230	230	212	210	207	201	197	194
TONAL + CPR$_0$ 90	M	49	49	47	47	47	47	46	44
	R	238	239	225	224	221	219	214	210
TONAL + CPR$_0$ 135	M	49	49	47	48	48	46	46	44
	R	232	232	217	216	215	212	209	203

Table 9.17 The influence of the texture displacement vector, tdv, and the four directions of **CLUSTER PROMINENCE** on the WOOD training set. lfn = 5.0; dz = 10.0; step = 0.002; loss ratio = 2.0; 2000 iterations; CPR coefficient = 50.0.

The most significant result to emerge from tables 9.11 to 9.17 is that in every measure the most significant texture information lies in the 0° direction and is discernable by the improvement in both M and R values over those for tonal measures alone. For the purpose of comparison the definition of a 'significant' improvement is one that misclassifies less than 44 local areas out of the 1073 of the training set. This is chosen because the best that is achieved with the tonal measures is 45/1073 using the mean and kurtosis (equation 9.1i). Using this definition the only texture measure that does not show a significant improvement is ASM which in the 0° direction reaches final values of M comparable to that achieved with tonal measures.

Entropy achieves a significant improvement for only the two values of tdv = 1 and 12 and is otherwise comparable to the tonal measures. Both ENT and ASM, for directions other than 0°, have values of M and R that are considerably higher than those for the tonal measures. Examination of the decision boundaries at the end of the iteration reveals that for all these instances the decision boundary has not moved much from the initial position. This indicates that the gradient at the initial decision boundary position is small. A more optimal position of the decision boundary would be reached by using a larger step in the iteration but the consideration of the results of the other texture measures suggest that no significant improvement over the tonal measures will result.

CON has a value of M of 40-41 in the 0° direction for tdv = 1 to 8 with a rise of M up to tonal levels for tdv = 12 and 16 and for all values of tdv in the other directions. This is consistent with the perception that, in the 0° direction, there is a high probability that pixels separated by a distance of tdv will be of similar grey level if the local area contains horizontal growth rings. This will give the S_0 SGLDM a low value of contrast as the elements of the matrix are clustered about the diagonal. (Recall that contrast is the moment of inertia of the SGLDM about the diagonal.) Feature areas will have a lower

probability that pixels separated by tdv will be of similar grey level; the elements of S_0 will be less clustered about the diagonal; and the value of the contrast will be higher.

Directions other than 0° lack spatial dependence because of the variation in the spacing of the growth rings. The texture coefficient of the decision boundary tends toward zero with iteration of the classification algorithm confirming that no useful information is present in these measures. This is further confirmed by the final values of M and R which are close to those for the tonal measures alone.

An examination of the differences between the local areas misclassified with the use of tonal measures and those misclassified with the use of tonal measures and a texture measure reveals the type of areas that the texture measures aid in detecting. Table 9.18 lists these differences for one example each of CON, COR, LHOM, and CSH for values of the tdv that result in the lowest value of M and R. It is found that CSH and CPR with a tdv = 6 misclassify exactly the same areas.

IMAGE	AREA	TONAL MEASURES	+ CON tdv=4	+ COR tdv=4	+ LHOM tdv=3	+ CSH tdv=6
WOOD01	3,4	FN	TP	TP	TP	TP
WOOD01	6,4	TN		FP		
WOOD05	0,5	FP			TN	
WOOD06	4,3	FP	TN	TN	TN	TN
WOOD09	5,1	FP	TN	TN	TN	TN
WOOD09	6,2	FP	TN	TN	TN	TN
WOOD10	3,4	TN	FP	FP	FP	FP
WOOD10	4,5	FP	TN	TN	TN	TN
WOOD10	5,2	FP	TN	TN	TN	TN
WOOD10	5,4	TN	FP	FP	FP	
WOOD12	5,2	TN	FP	FP	FP	
WOOD13	3,2	FN	TP	TP		
WOOD13	3,3	FN	TP	TP	TP	
WOOD14	3,4	TN	FP	FP	FP	
WOOD14	4,0	FN	TP	TP		
WOOD14	4,2	FN		TP	TP	TP
WOOD15	1,2	FP	TN	TN		TN
WOOD15	2,3	FP				TN
WOOD17	2,5	TP			FN	
WOOD17	3,3	FN	TP	TP	TP	
WOOD21	1,3	FN	TP	TP	TP	
WOOD21	2,2	FN				TP
WOOD22	6,0	TN	FP	FP		
WOOD22	6,2	TN			FP	
WOOD24	2,2	FN	TP	TP	TP	TP
WOOD24	2,5	TP		FN		
WOOD32	4,2	TN				FP
WOOD38	0,1	FP				TN
WOOD38	3,6	FN			TP	

Table 9.18 Differences in classification between tonal measures alone (equation 9.2) and with a single texture measure (0° direction) added. The differences are indicated by TN (true negative), TP (true positive), FN (false negative), and FP (false positive).

Referring to table 9.18, 7 areas that are misclassified with the tonal measures are correctly classified with the aid of any of the five texture measures. Two of these (WOOD01(3,4), WOOD24(2,2)) are the centres of large knots that are quite faint and are falsely classed as negatives with the tonal measures. The other 5 are correctly classed as negatives with the addition of texture information. All 5 areas contain only growth rings or dark resin: WOOD06(3,4), WOOD09(5,1), WOOD09(6,2) containing

horizontal growth rings or resin; and WOOD10(4,5), WOOD10(5,2) with the growth rings at an angle.

Of the other 12 areas that are correctly classed with the aid of one of the texture measures it is the deviation from horizontal growth rings, such as occurs with the centre of a faint intergrown knot, that enables the texture measure to identify the area as a feature local area. Conversely, it is the presence of horizontal growth rings (or rather the absence of vertical texture as in the case of a uniform area with no growth rings) that enables the texture measure to identify the area as a clear local area.

CSH and CPR both contain a large amount of useful texture information. This is noted by the result that the addition of either measure corrects 12 misclassifications of the tonal measures and only introduces 2 new misclassifications. Both of these (WOOD10(3,4), WOOD32(4,2)) are areas that contain a dark vertical growth ring.

CON, COR and LHOM correct 13 or 14 of the misclassifications of the tonal measures but introduce 5 to 7 new misclassifications. These are largely false positives due to resin stain (WOOD14(3,4)), or a dark growth ring at an angle (WOOD10(5,4)). LHOM misclassifies WOOD17(2,5) as a negative which results in a small knot being undetected. One false positive classed by CON, COR and LHOM is WOOD12(5,2) which is the edge of a large faint intergrown knot. This can be considered to be a true positive and reflects the difficulty in specifying a correct class for all the elements of the training set.

WOOD13 contains a faint intergrown knot that lies in four local areas ((2,2), (2,3), (3,2), (3,3)). All four areas are misclassified using tonal measures but two of the areas are correctly classed with the addition of CON or COR. LHOM corrects only one of the areas but CSH and CPR do not improve upon the tonal measures for this feature.

9.4 CONCLUSIONS OF THE USE OF TONAL AND TEXTURE MEASURES TO PERFORM THE CLASSIFICATION OF LOCAL AREAS

Analysing the results of the class/feature decision boundary requires some elaboration to clarify the meaning of the numbers. An important point to be made is that the number of misclassifications reached by the hyperplane is not the number of knots that go undetected. Some of the false negatives are local areas that contain a small portion of a feature, the rest of which is classified correctly. These areas are referred to as connected false negatives and the information contained within them may be retrieved by later processing. In this way only the unconnected false negatives are considered to contribute to error in a non-recoverable way.

The conclusion from results obtained with the addition of texture measures to tonal measures are:

a) The calculation of the texture measures represents a considerable computational penalty involving the calculation of an EPQ histogram and a SGLDM. The results show, however, that only the SGLDM in the 0° direction is required to obtain a significantly improved result.

b) The addition of texture measures to the tonal measures reduces the number of misclassifications from $M = 48$ to $M = 38$ (with cluster shade and cluster prominence, tdv = 6) which is a significant improvement of more than 20%.

c) No significant improvement was found with the use of angular second moment or entropy texture measures.

d) Contrast, correlation, local homogeniety, cluster shade and cluster prominence texture measures all contain information that improves the classification of feature local areas compared to the use of tonal measures alone.

The true effect of this improvement in the detection of feature local areas depends on how this information is used to discriminate the features. This is the subject of the next chapter.

10. FEATURE DISCRIMINATION

Up to this point this study has focused upon the detection of features. The images of radiata pine have been divided into non-overlapping subimages called local areas and the statistics of the grey level histogram have transformed each local area into a sample point in multi-dimensional pattern space. A linear decision boundary, derived from a set of labelled training local areas, has been used to describe a hyperplane in pattern space that classifies any introduced sample point as either a feature local area or a clear local area. This is the feature detection process that is described in detail in the preceding chapters. It enables the clear areas to be identified and removed from further processing.

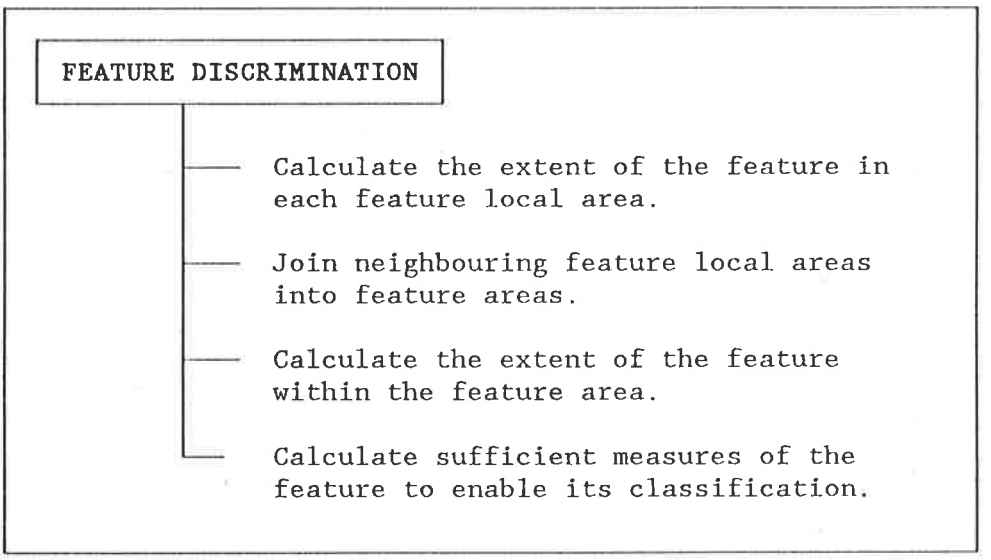

Figure 10.1 The outline of feature discrimination once the feature local areas have been detected.

The feature local areas that remain can be one of several different types:

(a) The area can be a false positive, that is, a clear area that has been wrongly classified as a feature local area.
(b) The area can wholly contain a feature in which case the feature must be identified.
(c) The area can contain part of a feature with the rest of the feature in adjoining local areas. These neighbouring feature local areas must be merged to form a feature area that encompasses the entire feature which is then identified. The proposed process is illustrated in figure 10.1.

Within a feature area there may be a number of separate features which must be discriminated. The dimensions and tonal measures of each feature are postulated as being sufficient to classify the features correctly. The methods that are used to obtain these measures are described in this section.

10.1 DETERMINING FEATURE EXTENT WITHIN A LOCAL AREA

Before the neighbouring feature local areas can be connected to form larger feature areas, a method is required to determine the extent of the feature within each local area. Only if the portion of the feature within a local area extends to the border can it be considered that the feature may also be contained within the neighbouring local area.

Ideally, the technique that determines the extent of a feature within a local area will also detect if the area is a false positive. If false positives are not detected at this stage then they must be detected at a later stage after more processing. Two different methods are examined for their suitability in detecting the extent of the feature within a local area, and the effectiveness of each method is assessed.

In the diagrams that follow the local area numbering scheme illustrated in figure 10.2 is used.

	0	1	2	3	4	5	6	X
0	0,0	1,0	2,0	3,0	4,0	5,0	6,0	
1	0,1	1,1	2,1	3,1	4,1	5,1	6,1	
2	0,2	1,2	2,2	3,2	4,2	5,2	6,2	
3	0,3	1,3	2,3	3,3	4,3	5,3	6,3	
4	0,4	1,4	2,4	3,4	4,4	5,4	6,4	
5	0,5	1,5	2,5	3,5	4,5	5,5	6,5	
6	0,6	1,6	2,6	3,6	4,6	5,6	6,6	
Y								

Figure 10.2 Numbering of local areas in an image.

10.1.1 THE LINE STATISTICS METHOD

This method relies on the fact that a feature is significantly "different" from the background of growth rings and that this difference is detectible within each row and column of pixels in the local area. The aim is to place a box around the feature within the local area to define its extent.

Intuitively, and from the experience with the tonal measures of the linear decision boundary, the mean and the variance should provide good measures for finding the difference that is the edge of a feature. Starting from the edge of the local area and working inwards the mean and variance are calculated for each row, from top-to-bottom and from bottom-to top, and each column, from right-to-left and from left-to-right. If the mean value of the line of pixels is below a set threshold or the variance is above a set threshold then that line is marked as the extremity of the feature. The process continues until the horizontal and vertical extremities are defined.

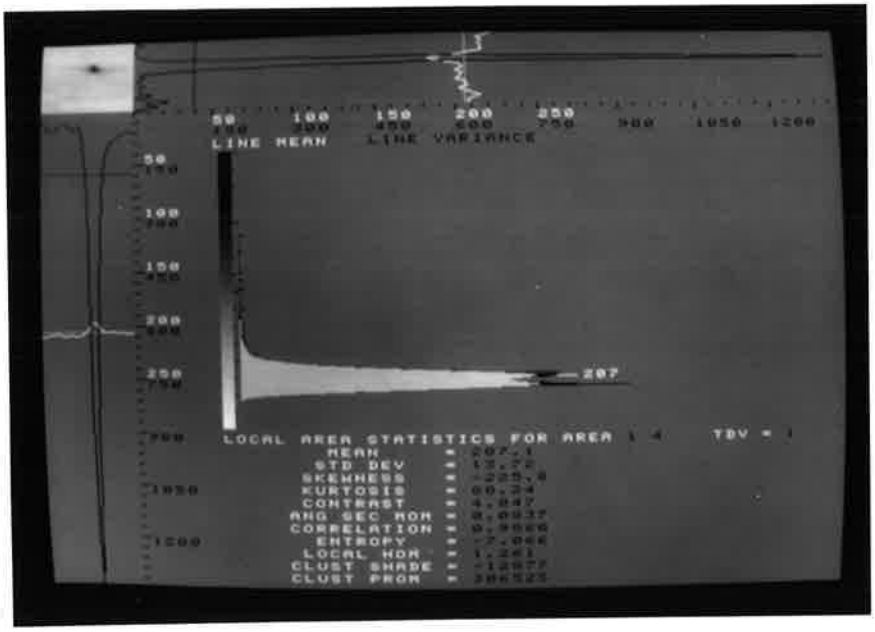

Figure 10.3 The statistics of local area (1,4) of WOOD42.

The ideal situation can be appreciated with reference to figure 10.3 which displays the statistics for local area (1,4) of WOOD42. A small pin knot, actually an element of needle trace, gives the local area a position in tonal pattern space that classifies it as a feature local area. The scale to the right of the local area gives the value of the mean and variance of each horizontal row of pixels in the area. Similarly, the scale below the local area gives the value of the mean and variance of each vertical column of pixels.

The influence of the feature on the line statistics is clear. Both the row and column variance values are significantly increased for those rows and columns that contain the feature. Because the feature occupies only a small portion of a row or column the influence on the mean values is small but noticable. A threshold value of the line variance is sufficient to define the extent of this feature.

Figure 10.4 The statistics of local area (2,0) of WOOD05.

A local area that contains a portion of pith along its bottom edge is shown in figure 10.4. The edge of the pith is not quite horizontal and the line variance exhibits a large increase in the rows that contain this transition from pith to background wood. The centre part of the pith, however, is relatively uniform and so has a low line variance. The line variance alone can only distinguish the edge transition and if the pith had been precisely horizontal it would not have been capable of detecting even this. However, the line mean of the pith is significantly lower than that of the background wood and so a line mean threshold can successfully be used to identify the body of the pith as feature area.

The line statistics method combines the line mean and the line variance in a logical threshold structure of the form:

IF (line variance > line variance threshold)
OR (line mean < line mean threshold)
THEN this row/column marks the edge of the feature.

This process works well for well defined features such as a dark pin knot on a light background. It does not work so well for a light knot or a feature with dark growth rings surrounding it. Patches of bark or resinous growth rings extend the feature area making impossible accurate sizing of the feature extent. This over-sizing of the feature extent can be illustrated with reference to figure 10.5 which shows the feature local areas of image WOOD01 and the extent determined with a variance threshold of 250 and a mean threshold of 165. The extent of the small hole is extended by the variance due to specular reflections. This can be seen in the statistics of the local area (1,2) in figure 10.6.

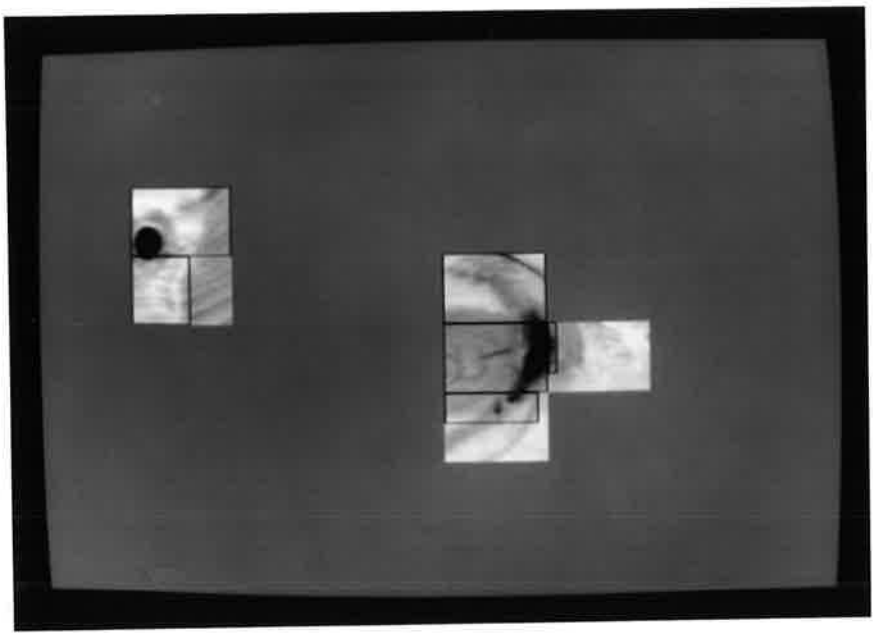

Figure 10.5 The extent of feature local areas of WOOD01 using the line statistics method.

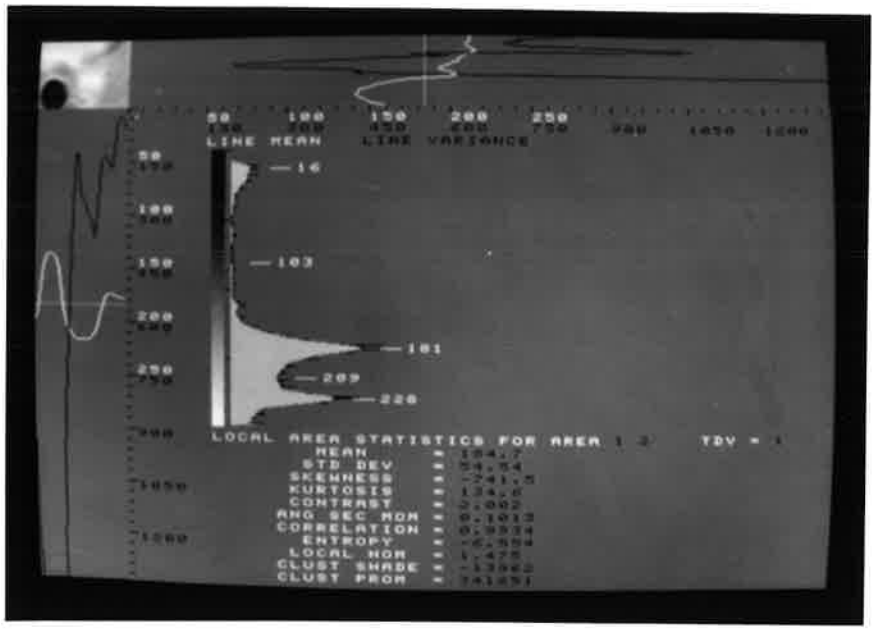

Figure 10.6 The statistics of local area (1,2) of WOOD01.

The over-sizing of the local areas of the large knot is due to the dark growth rings that merge with the knot. This method of determining the feature extent does, however, confirm the presence of a feature within the local areas and provides important connectivity data.

The greatest problem with this method is the setting of mean and variance threshold limits. Using a specially written interactive software package it is found that the most suitable values for the mean and variance thresholds vary over a range. The value of the line variance threshold is satisfactory over quite a large range, with large threshold values (>500) resulting in the more smoothly varying features being defined small in extent, and small threshold values (<150) resulting in feature extent being defined to the edge of the local area. This latter result is undesirable because a false indication of connectivity is given. Choosing a variance threshold between these two limits gives good results. A value of 250 is found to provide a balance between the feature within the local area being defined too large and too small.

The most suitable value for the mean threshold is complicated by the fact that the overall brightness of the image of wood is variable, being a function of the lighting, the type of wood (heartwood or sapwood), and the gain of the camera. The difference in grey level between pith and background is very large (100 grey levels, from figure 10.4) so the choice of a mean threshold solely to detect the edge of pith is satisfactory between the values of 80 and 170. The use of the mean threshold to detect the edge of smoothly varying features such as spike knots requires a threshold toward the higher end of this range (150-170). A value of 165 is used in the results presented in Appendix 2A and 2B.

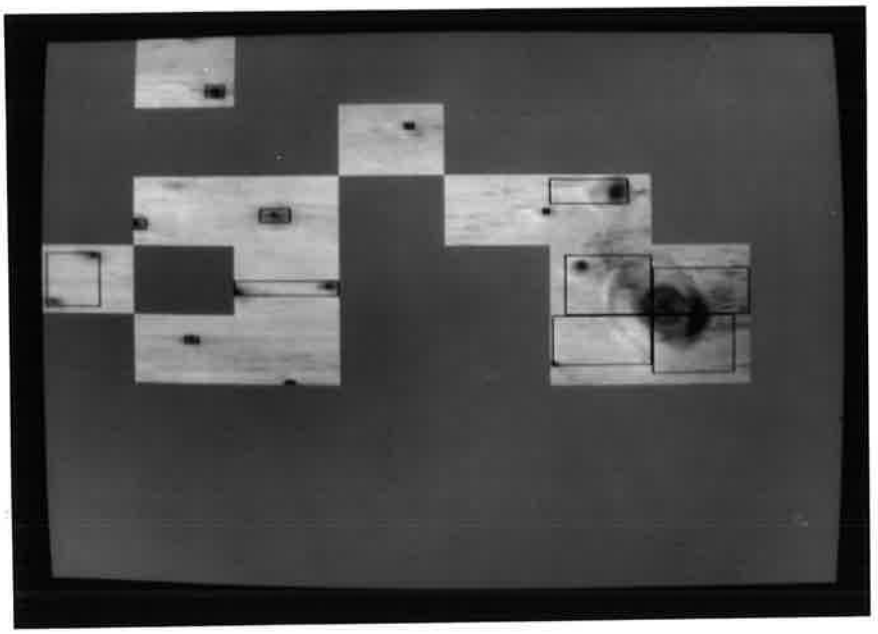

Figure 10.7 The extent of feature local areas of WOOD42 using the line statistics method.

A difficulty of definition exists with this method. Figure 10.7 shows the situation where some feature local areas contain more than one feature, for example, (0,3), (2,3) and (5,3). This method contains no information that allows the separate features to be identified individually.

A further difficulty is the detection of false positive areas. Table 10.1 lists the minimum variance threshold and the

maximum mean threshold needed to enable the false positive areas of the training set, as classified by the decision boundary of equation 9.2, to be classed as clear.

FALSE POSITIVE LOCAL AREA	MIN. LINE VARIANCE THRESHOLD	MAX. LINE MEAN THRESHOLD
WOOD04 (2,3)	>1350	160
WOOD05 (0,5)	>1350	150
WOOD06 (2,4)	1100	155
(4,3)	1200	170
(5,5)	1080	133
WOOD09 (4,1)	660	110
(5,1)	230	119
(6,1)	575	100
(6,2)	390	110
WOOD10 (2,4)	>1350	160
(4,5)	1200	160
(5,2)	1080	155
(5,5)	1230	132
(6,3)	580	150
WOOD12 (2,2)	900	150
WOOD15 (1,2)	1020	172
(1,5)	1200	164
(2,3)	930	150
WOOD22 (2,3)	1100	160
WOOD38 (0,1)	960	143

Table 10.1 Line thresholds required to classify false positive local areas as clear. False positives defined by the decision boundary of equation 9.2.

The data in table 10.1 indicates that the line statistics can not be used to recognise false positive areas. Increasing the variance threshold or lowering the mean threshold in order to detect a small number of these false positives leads to degradation of the method's ability to determine the feature extent in true positive areas.

10.1.1.2 CONCLUSIONS OF THE LINE STATISTICS METHOD

The line statistics method is able to detect the presence and extent of features within feature local areas. Its limitations are:

a) Its reliance upon fixed threshold values. The variance threshold is independent of the overall brightness level of the image but the mean threshold is very dependent on this. It is desirable for an industrial system to be tolerant of lighting variations and robust in the event of expected changes. The line statistics method, with its reliance on a fixed mean threshold, does not possess this robustness.

b) There is also no capability of separating individual features if there is more than one within a local area. This method describes the limits of the features within a local area as that of one feature. There exist instances in the sample set of images where a local area contains several features or portions of features. The line statistics method is not suitable for the discrimination of separate features.

c) The method is unable to detect false positive local areas. These are passed onto the next stage of processing.

10.1.2 THE ADAPTIVE THRESHOLD METHOD

This method relies on the fact that features are most often darker than the background wood. By the choice of a suitable threshold value the local area can be binarised, that is reduced to only two grey levels (black and white). This creates an image with the feature area black and the background white which has the advantage of being a simpler representation of the feature.

The difficulty is to find the threshold value which best binarises the local area into feature and background. The source that supplies this information is the grey level histogram which is constructed at the feature detection stage. A dark feature on a light background that comprises from 5% to 95% of the local area will have a bimodal histogram with two peaks, one corresponding to the feature and the other to the background. The choice of a threshold value is dependent on the identification of the peaks and valleys of the histogram. Choosing a threshold in this way isolates the choice from global variations and makes it responsive to local influences. The threshold choice is adaptive

being dependent upon the statistics of the histogram of each local area.

10.1.2.1 FINDING PEAKS AND VALLEYS IN THE HISTOGRAM

The histogram data is always variable and requires smoothing. A triangular seven element mask filter is used of the form:

$$n'_g = (n_{g-3}+3.n_{g-2}+5.n_{g-1}+7.n_g+5.n_{g+1}+3.n_{g+2}+n_{g+3})/25$$

where n_g is the grey level. This is applied twice to the histogram array to provide a smooth curve for the peakfinding routine.

The peakfinding routine starts from the low intensity end of the histogram and starts looking for a peak. It lists the location (grey level) and size (number of pixels) of each peak and valley in the histogram using the method described in figure 10.8. The sensitivity of the peakfinding routine is adjusted by altering two parameters. The breadth (5 in figure 10.8) is the number of grey levels that must be incremented before a peak (valley) is confirmed and the search switched to that for a valley (peak). This has the effect of suppressing very small peaks and valleys that are close together. To confirm a peak (valley) the breadth criterion states that the number of pixels at the current grey level must also be 4 less than (greater than) the value at the peak (valley). This has the effect of suppressing small peaks and valleys that are isolated.

Together the breadth and depth values determine the sensitivity of the peakfinding routine. All the local areas are of the same size and the values of 5 for the breadth and 4 for the depth are found to give satisfactory results. Smaller values result in the identification of many more peaks and valleys and higher values result in significant valleys and peaks being overlooked. As is the case in many engineering decisions the choice of values reflects a compromise that is considered optimal based upon the available information.

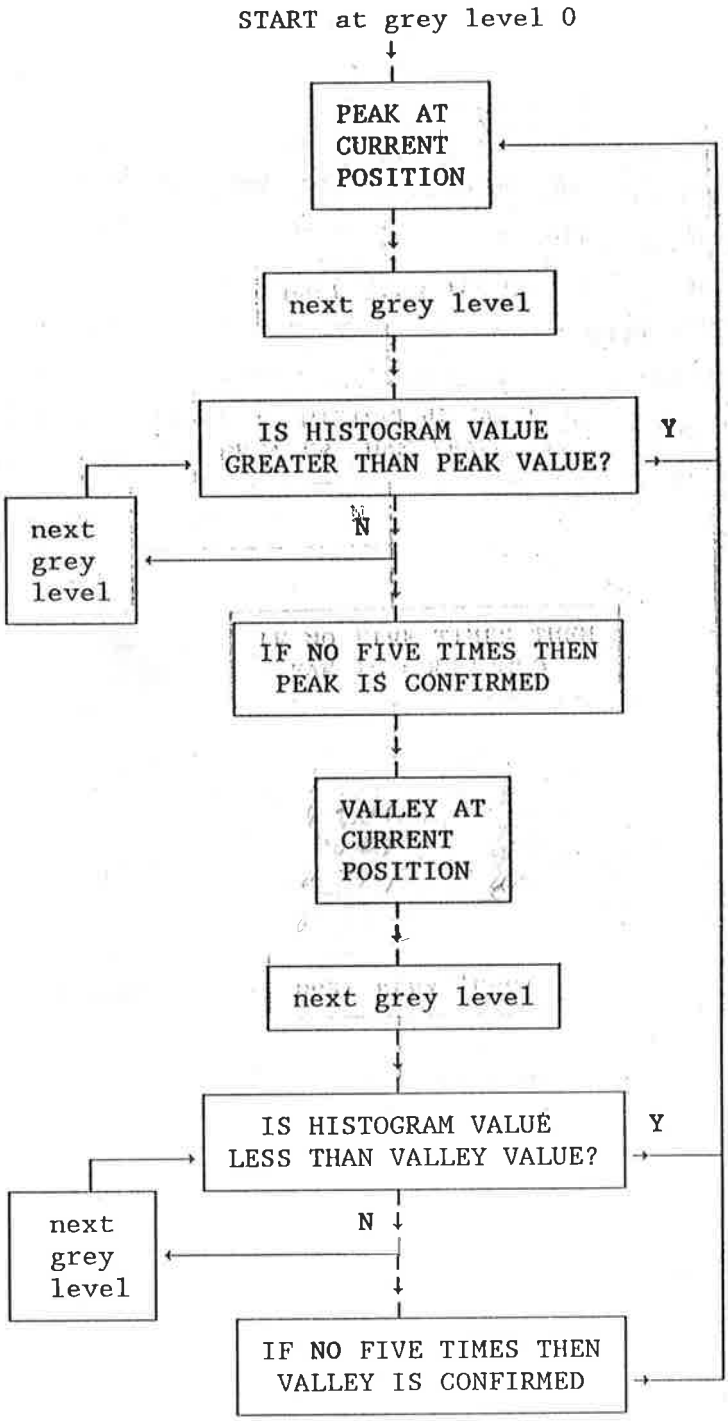

Figure 10.8 Flow diagram for finding peaks and valleys in the grey level histogram.

10.1.2.2 CHOOSING A THRESHOLD VALUE

Ideally the feature local area has a bi-modal histogram with two peaks and one valley between them. The peakfinding routine finds this valley and the area is thresholded into a feature and background binary image. Unfortunately the variable nature of the wood results in a significant number of feature local areas having something other than an ideal bimodal histogram. The method developed to interpret the grey level histogram and choose a satisfactory threshold value is illustrated in the flow chart of figure 10.9. The reader is urged to refer to this figure in the discussion of the method that follows.

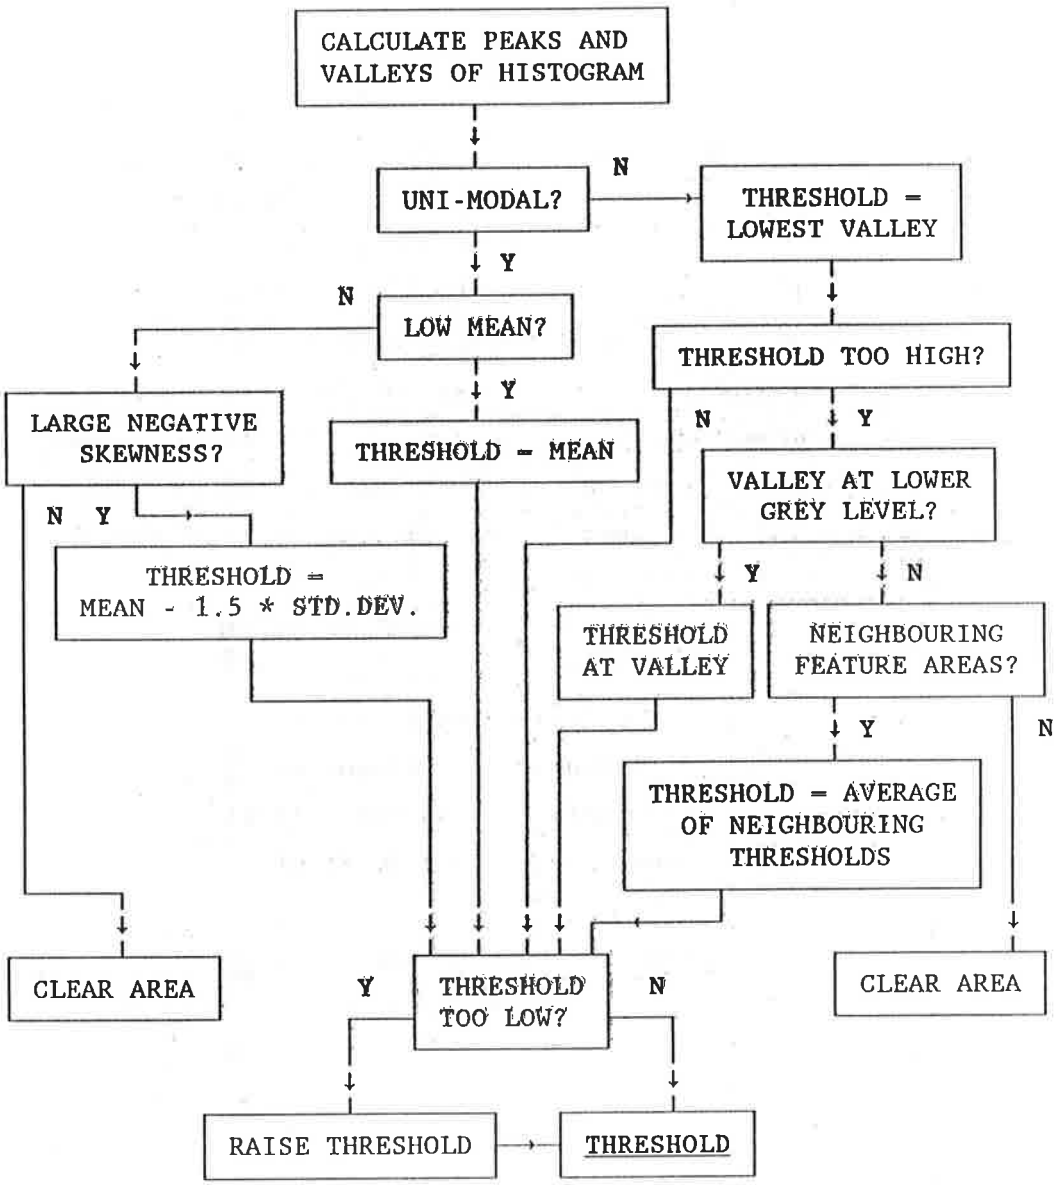

Figure 10.9 Flow diagram for determining the threshold value for fixed size feature local areas based on the grey level histogram.

The linear decision boundary classifies local areas as feature local areas based upon the tonal statistics of the grey level histogram. For the most part clear areas are characterised by a **uni-modal histogram** that has a relatively high mean, a relatively low standard deviation and kurtosis, and a skewness close to zero. An example is the local area (6,1) of the image

WOOD01 (figure 10.10b). The difference between the grey levels of the light and dark areas of the growth rings is relatively small and the transition between the two is smooth with many pixels having intermediate values.

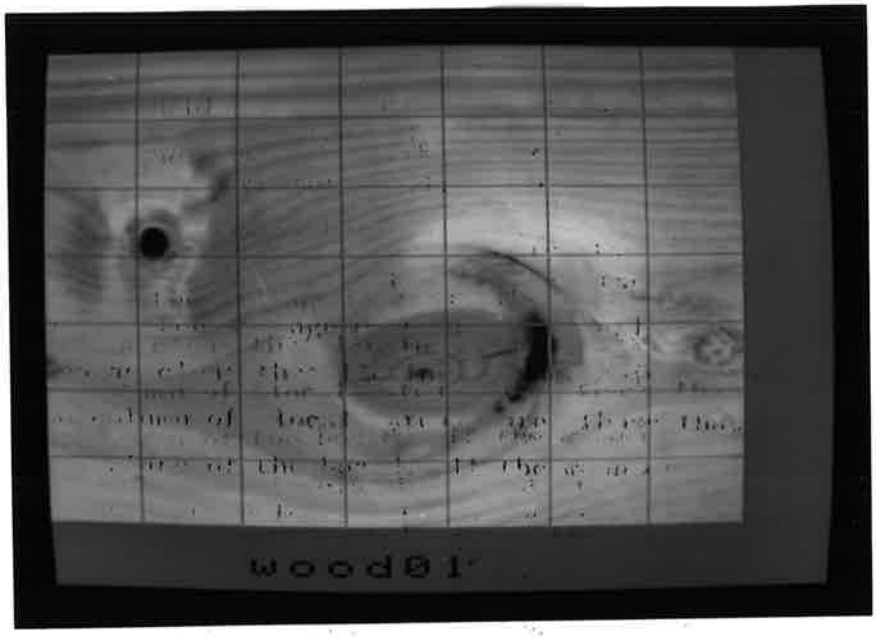

Figure 10.10a The division of WOOD01 into local areas.

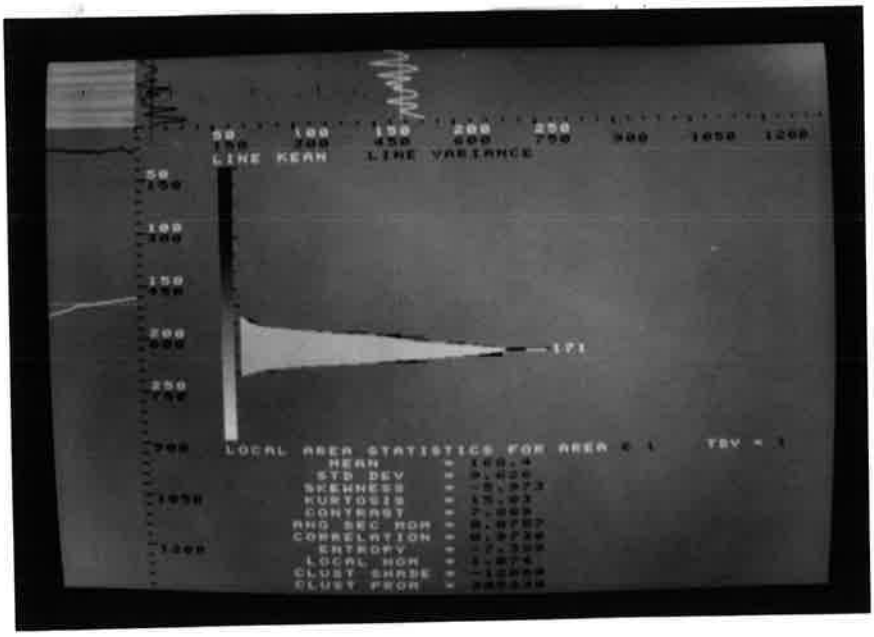

Figure 10.10b The statistics of local area (6,1) of image WOOD01.

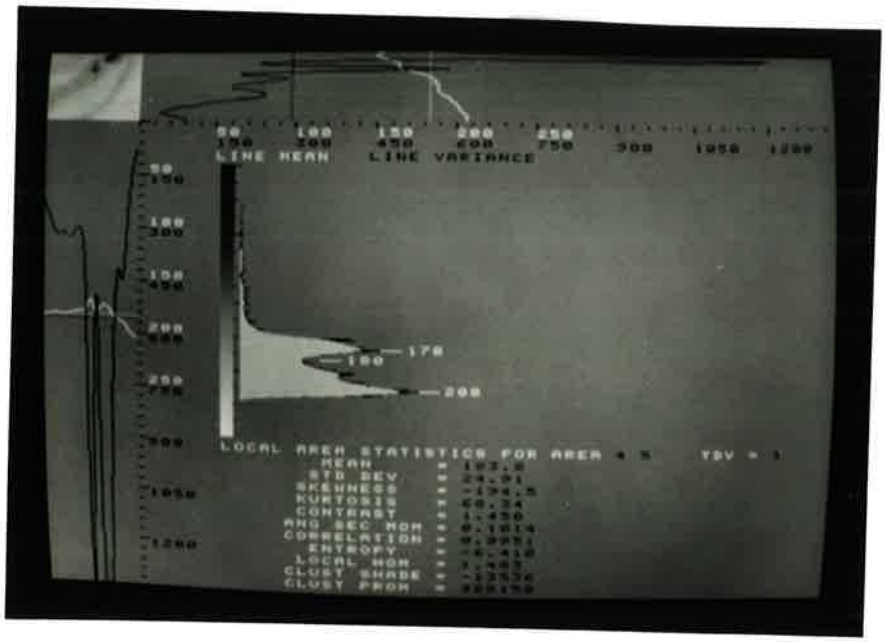

Figure 10.10c The statistics of local area (4,5) of image WOOD01.

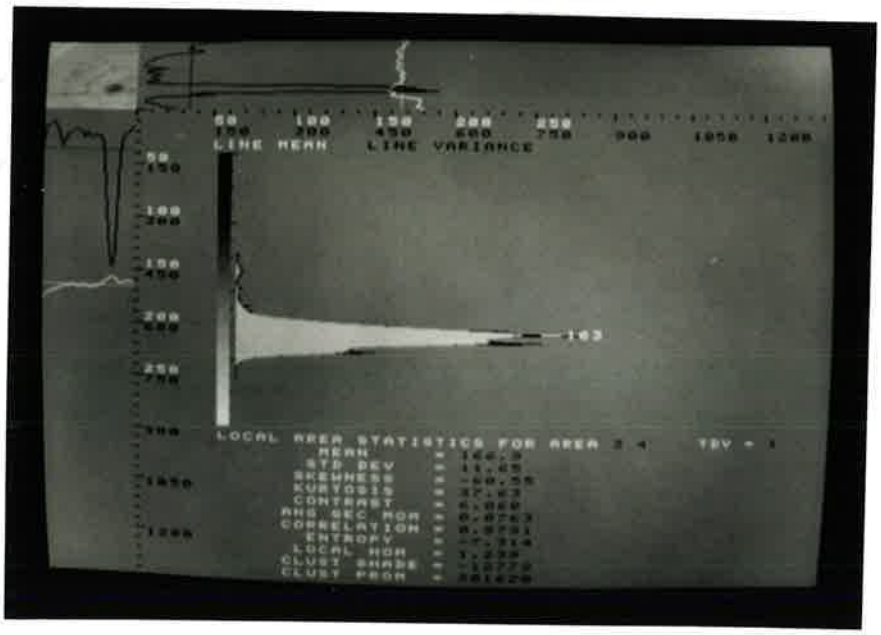

Figure 10.10d The statistics of local area (3,4) of image WOOD01.

The peakfinding routine described in section 10.1.2.1 finds that 44 of the 239 feature local areas of the training set have only a single peak. There are three reasons for this. The

first is that the area lies fully over a dark feature such as the centre of a dark knot. In this case the area will have a low mean value and the threshold can be set in relation to the mean. But what is a low mean? Generally it is best to use measures that are independent of the mean which is known to vary from plank to plank and even within a plank. The use of the mean grey level becomes unavoidable in the discrimination of wholly dark regions because no other measure will suffice. The method chosen to distinguish an area as being sufficiently dark to be recognised as the centre of knot is to compare the mean of the area in question with the average mean of the local areas that have been classified as clear that lie in the same column, the column clear mean. The column of local areas are those that form a section across the width of the board. If the mean of the suspected area is less than 30 grey levels below the column clear mean the area is classed as a feature local area and thresholded at its mean grey level. This is sometimes somewhat low but is satisfactory to provide connectivity with neighbouring feature local areas. This situation arises in 11 of the 44 uni-modal areas of the training set.

The second cause of a uni-modal histogram is the situation of a largely clear area with a very small part of a dark feature of which the local area (0,1) of the image WOOD21 is an example (figure 10.12b). The small number of darker pixels corresponding to the knot are seen as a small trail to one side of the peak. This gives the histogram a large negative skewness as well as a large kurtosis. The peak-finding routine lacks the sensitivity to find a valley in the tail. A threshold can be selected by choosing a value that is between one and two standard deviations less than the mean value. A value of:

threshold = mean - 1.5 * s.d.

is found to be satisfactory in these situations which number 28 of the 44 uni-modal areas of the training set. This rule is used if the skewness of the local area is less than -60.

The third cause of a uni-modal histogram are areas that contain dark growth rings or the edge of a faint knot or diffuse

spike knot. If the area does not trigger the rules for a low mean or a large negative skewness, it can be classified as clear. This occurs with 5 of the 44 uni-modal areas of the training set. Examination of these 5 areas reveals that they are all true positive feature local areas.

The area (1,5) of WOOD38 contains the diffuse edge of a large spike knot and the exclusion of this area does not affect the definition of the knot area. The areas (3,1) and (3,2) of WOOD15 contain the edge of a very faint knot. A threshold that defines the extent of this feature is impossible to find because the feature is so faint. The area (3,5) of WOOD15 contains a small knot with an indistinct boundary. The neighbouring local area that contains the other half of this knot has a multi-modal histogram for which a threshold is not found. The result is that this knot is classed as clear even though both of the local areas that contain part of it are classed correctly as feature local areas. The area (2,5) of WOOD19 contains half of a medium knot, the other half of which lies in area (1,5).

This suggests that, rather than classifying these areas as clear, the neighbouring local areas should be examined and the threshold set to the average of the thresholds of any neighbouring feature local areas. This would not affect WOOD38 (1,5) which does not affect the definition of the feature area, nor would it benefit WOOD15 (3,1) and (3,2) apart from increasing the extent of the feature area because of the faint nature of this knot. The neighbouring feature local area of WOOD19 (3,5), that is (3,6), has an undefined threshold and so cannot be used to define the threshold of (3,5). A satisfactory threshold can be applied to WOOD19 (2,5) by using the threshold of (1,5) and leads to a more complete definition of the feature.

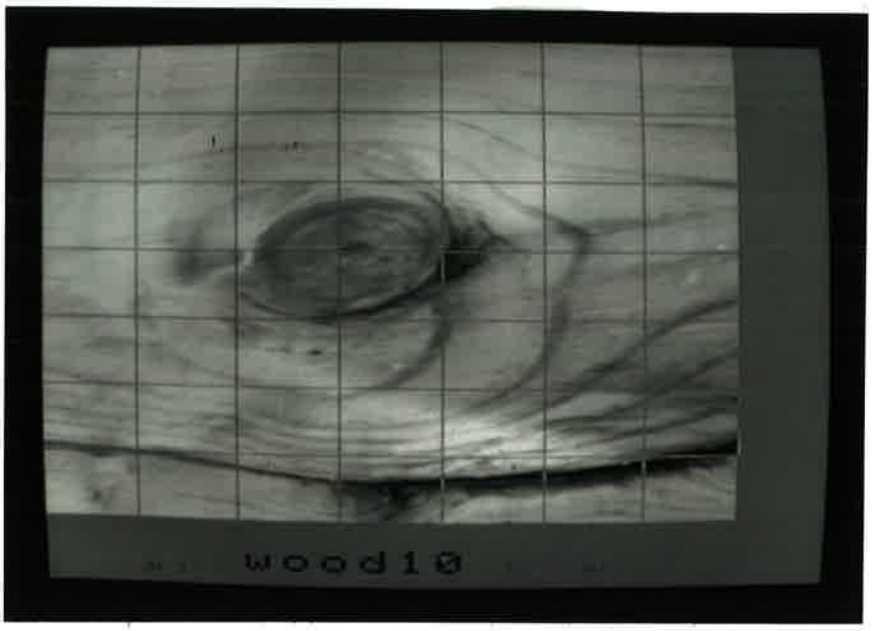

Figure 10.11a The division of WOOD10 into local areas.

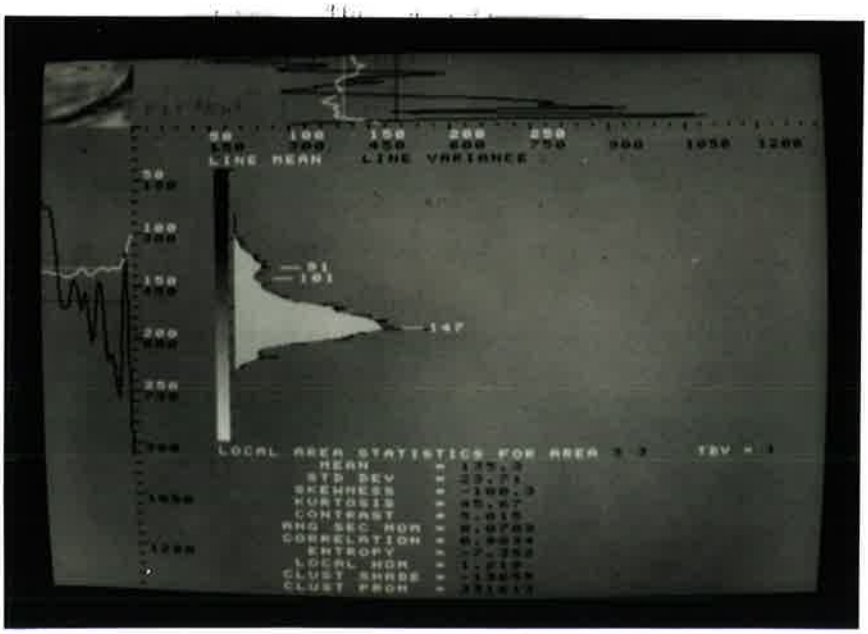

Figure 10.11b The statistics of local area (3,3) of image WOOD10.

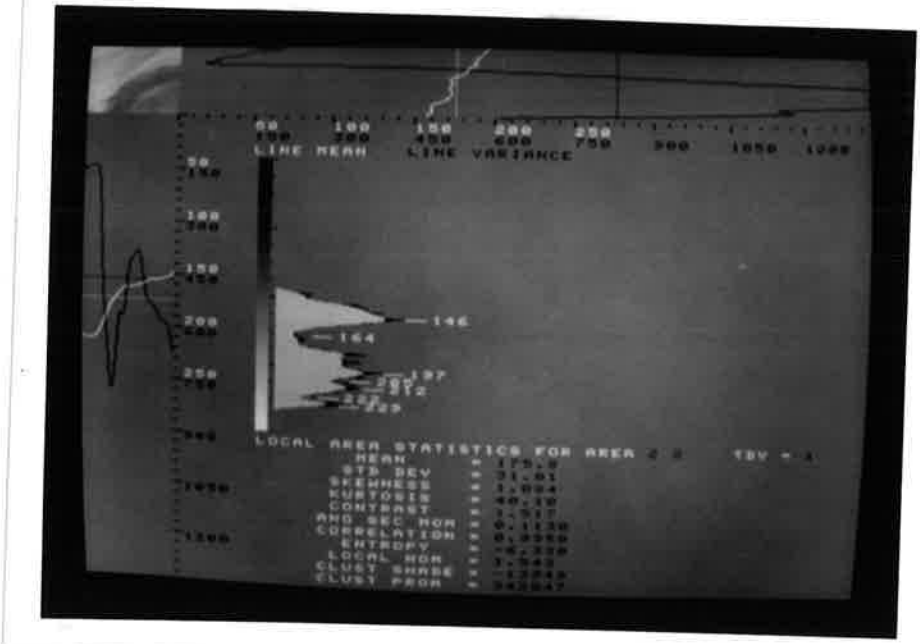

Figure 10.11c The statistics of local area (2,2) of image WOOD10.

The remaining 187 feature local areas have bi-modal or multi-modal histograms. The local area (2,2) of the image WOOD21 (Figure 10.12c) shows a clear area with a small portion of a larger knot. This is an example of an ideal histogram that is clearly bimodal with the higher intensity grey level peak representing the lighter background and the lower intensity grey level peak representing the dark feature area. The area under each peak represents the proportion of the local area covered by that part of the histogram. In these cases the optimum threshold is the in the lowest valley between the peaks.

Nearly half of the local areas that are classed as feature areas contain three or more peaks in the histogram. The local area (2,2) of image WOOD10 (figure 10.11c) is best thresholded at the grey level valley of 164, the higher valleys of 205 and 222 due to growth rings. The optimum threshold is found by choosing the valley with the lowest number of pixels, that is the lowest valley. This is because the depth of the valley between the peaks is an excellent indication of the separation of distinct areas of different average grey level.

Problems of interpretation arise when the two areas of feature and background are not distinct. The image WOOD21 is typical of spike knots in general in that the knot merges into the growth rings in a smooth diffuse manner. The local area (3,4) (figure 10.12d) is best thresholded at the lower grey level valley at 163 but the lowest valley rule places the threshold at the grey level of 185. This is too high and a conditional is placed upon the threshold at this stage. A threshold that is too high creates a feature area that is too large which in turn may falsely indicate connectivity to neighbouring local areas. A value of 10 grey levels less than the column clear mean is taken as the maximum value of a threshold. If the first threshold value is too high then the threshold is set to the value at the first valley below this upper limit. In the case of local area (3,4) of WOOD21 this upper limit is 172 and so the threshold is set to the valley at 163.

Figure 10.12a The division of WOOD21 into local areas.

Figure 10.12b The statistics of local area (0,1) of image WOOD21.

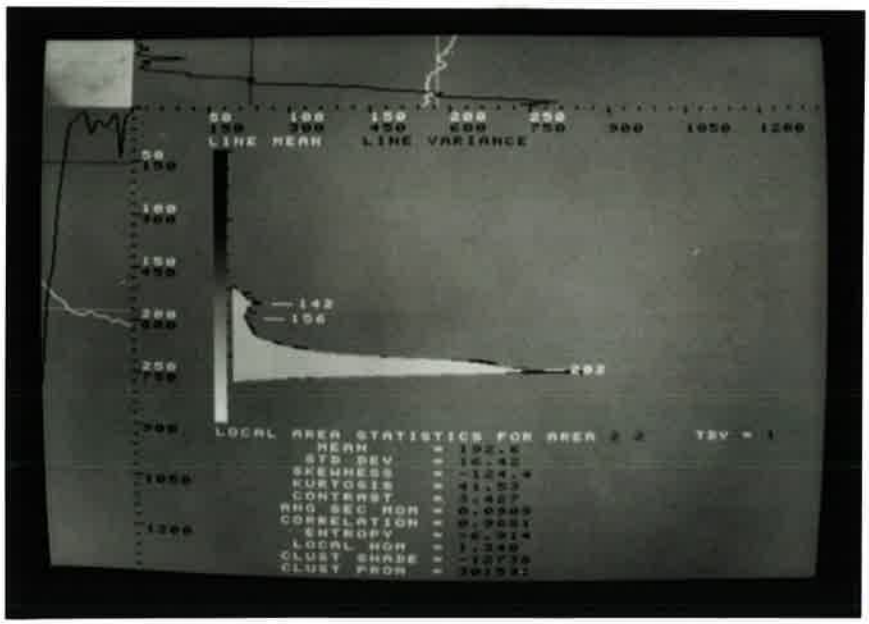

Figure 10.12c The statistics of local area (2,2) of image WOOD21.

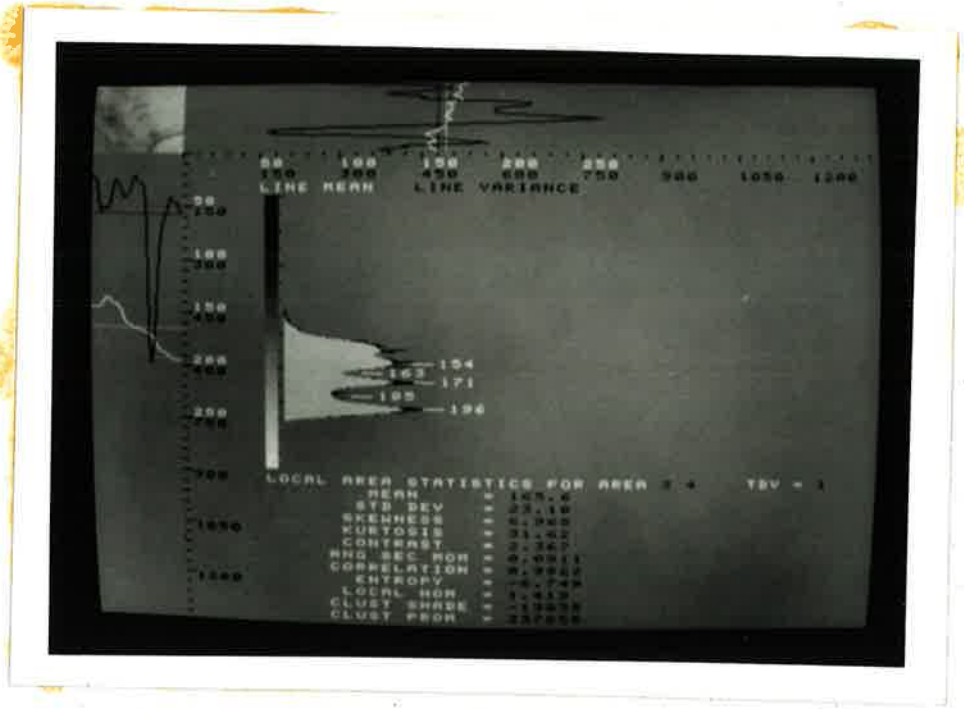

Figure 10.12d The statistics of local area (3,4) of image WOOD21.

If no such valley exists another method is used to determine a threshold. The four neighbouring local areas are examined and the threshold is set to the average value of the threshold calculated for those that are feature local areas. This provides a satisfactory threshold for 19 of the 20 feature areas that fall into this category. The one area remaining has no neighbouring feature local areas and is classified as clear. Examination of this area (local area (3,6) of WOOD15) reveals that it is a diffuse, faint knot half of which is in the neighbouring local area (3,5). This latter area is one of the four uni-modal areas that is classified as clear and so the entire knot is lost. The reason for deciding to label areas such as (3,6) of WOOD15 as clear when they have no neighbouring feature local areas is to eliminate isolated false positives due to growth rings. In fact, the 4 growth ring false positive areas at this stage all have neighbouring feature local areas and so are not eliminated by this tactic.

A bi-modal histogram does not always mean that the correct threshold is at the valley. A common error is the situation of a local area that contains a portion of bark and

knot. The local area (3,3) of image WOOD10 (figure 10.11b) shows that the value of 101 at the valley of the histogram is too low and divides the image into knot and bark when the ideal threshold would be such that the knot area would also be feature area. This error of a threshold value that is too low is also encountered with all the other derivations used and its significance is with the problem of connectivity. The aim of defining the feature area extent within each feature local area is to join neighbouring areas into larger feature areas that encompass the whole feature. This can only be done if the individual local areas are connected. A lower limit for the threshold is set at 70 grey levels less than the column clear mean.

Dark growth rings can often make an ideal threshold difficult to identify. The local area (4,5) of image WOOD01 (figure 10.10c) contains a small amount of dark bark which gives the histogram a tail. The growth ring and knot are of very similar grey level and the valley of the histogram separates them from the background. The growth ring has the effect of extending the feature area to a size much larger than the knot. If the knot is darker than the growth rings then the latter are often eliminated when the feature area formed from the connected local areas is thresholded but if not then the error remains.

The interior of the large knot of WOOD01 is shown in figure 10.10c (local area (3,4)). The histogram is unimodal with a small number of dark pixels trailing to one side in the manner of figure 10.12b. However, in this local area the dark pixels correspond to the pith of the knot and the correct threshold is on the lighter side of the peak. Although it is difficult to distinguish this type of local area from the example of figure 10.12b the real difficulty is that this particular area is classed as clear by the decision boundary derived from tonal measures, that is, it is a false negative. The knot is not of sufficient contrast to be classed as a feature. However, the bark that accompanies it is classified correctly and the knot can be inferred from the bark which gives a good indication of the knot size.

10.1.2.3 SUMMARY OF THE ADAPTIVE THRESHOLD METHOD

The calculation of the peaks and valleys of the histogram is dependent upon the values of the breadth and depth of the peakfinding routine. Considering the values for these that have been stated and examining the identified feature local areas leads to a number of conclusions:

a) Uni-modal histograms are characteristic of three types of local areas. The first are clear areas, that is false positives. The second are areas that have a small portion of a feature that is either too small to create a separate peak or blends in smoothly to the main peak. The third are areas such as the centres of large knots and spike knots which contain little if any background area.

b) Areas with a small portion of dark feature and a unimodal histogram always have a large negative skewness and this can be used very effectively to calculate a threshold value. Areas that do not have a large negative skewness or a low mean are identified as false positives and are eliminated from further processing.

c) The threshold of local areas with bi-modal and multi-modal histograms is taken as the valley with the lowest number of pixels as being the best separator of feature area from background area. Where this threshold is too high (greater than the column clear mean minus 10) then it is obvious that this threshold is unsatisfactory. In this situation the area is thresholded either at a valley lower that the column clear mean or at a value equal to the average threshold of neighbouring feature local areas.

d) A final check is made to correct threshold values that are too low which lead to an inadequate amount of the feature area being detected and can lead to a loss of connectivity with neighbouring feature local areas. A value of the column clear mean minus 70 is taken as the minimum threshold value and all lower values are increased to this.

The threshold determined by this method is not always ideal and often leads to a large value of the feature extent,

illustrated in figure 10.13 for the image WOOD01. This can be compared to the result for the same image using the line statistics method in figure 10.5.

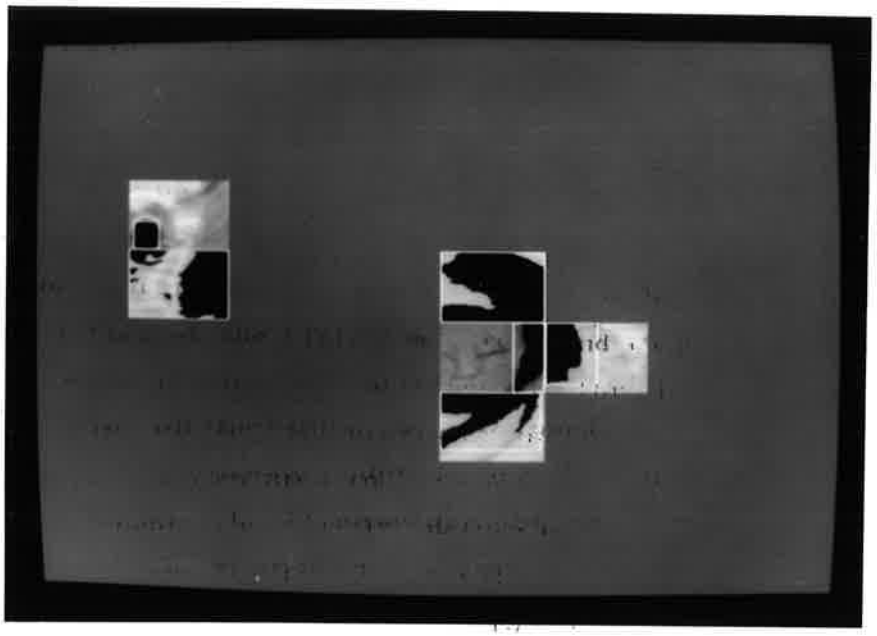

Figure 10.13 Determining the extent of feature local areas of WOOD01 using the adaptive threshold method.

Figure 10.14 shows the result of using the adaptive threshold method to determine the local area feature extent for the image WOOD42 (compare with figure 10.7). This also reveals that an ideal threshold is not always determined with this method, revealed by the large feature area in area (1,0). The figure does show, however, that this method is capable of discriminating the separate features within an area as separate blobs. This is not required at this stage in the connection of feature local areas but is a useful aspect that can be applied to the discrimination of features within a feature area.

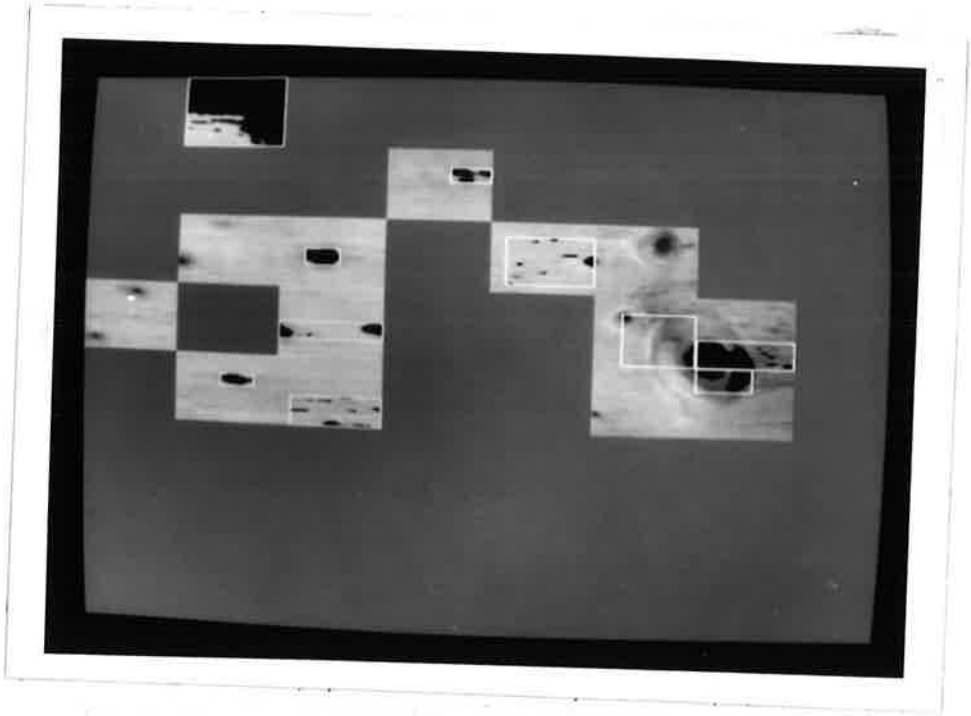

Figure 10.14 Determining the extent of feature local areas of WOOD42 using the adaptive threshold method.

10.1.2.4 CONCLUSIONS OF THE ADAPTIVE THRESHOLD METHOD

The adaptive threshold method proves satisfactory in defining the extent of features within feature local areas. The important points of the method are:

a) The thresholding of the area provides a means of discriminating the separate features within an area. This is not needed at this stage.

b) The calculation of the threshold is isolated from gross lighting levels. This gives the method a degree of robustness that is desirable in an industrial situation.

c) the method proves unable to distinguish false positives and falsely classifies a number of true positive areas as clear.

10.1.3 ASSESSMENT OF FEATURE EXTENT METHODS

The aim of detemining the feature extent within each feature local area is to establish connectivity, that is, to join together local areas that contain parts of the same feature into an encompassing feature area. Both methods developed fulfil this aim.

The line statistics method provides a satisfactory method for determining the extent of the feature area by using a logical combination of the threshold of the line mean and line variance. It is unsatisfactory for determining the number of features within a feature local area. It also lacks a degree of robustness because of its reliance on a fixed mean threshold.

The adaptive threshold method also provides a satisfactory method for determining the extent of the feature area by using the form of the grey level histogram to determine a threshold value with which to create a binary image. This segments each local area into feature and background on the assumption that feature area is darker than background area. The characteristics of the histogram are found to provide sufficient information to provide a satisfactory threshold in most cases. This has the advantage of providing information about the number of features within an area (not needed at this stage). It is also inherently robust, the adaptive nature of the algorithm isolating the threshold from the lighting level.

False positives prove not to be a serious problem. All the false positives are due to growth rings that are very dark and/or sharply contrasting or to dark resinous marks on the timber. Of the 20 false positives 14 do not affect the definition of the feature area apart from extending the feature area size marginally in 5 cases. In the remaining 6 cases the false positive areas form separate feature areas and threshold the growth ring into an unconnected feature area. These 6 are passed onto the next stage of processing for identification.

False negatives result in 3 faint intergrown knots passing undetected. A further 10 knots are defined only by the partial bark encasement that lies to one side of the knot. These cases of faint knots with partial encasement are due to bark becoming trapped in the angle of a steep branch. The knot is faint because it is young (as can be seen in the images of WOOD12, WOOD14, WOOD15, and WOOD17 in figures 5.2c, d, and e) and the bark length gives a good indication of the size of the knot.

The remaining false negatives do not affect the formation of the feature area or reduce its size only marginally.

10.2 FORMATION OF FEATURE AREAS FROM FEATURE LOCAL AREAS

The extent of the portion of the feature within each feature local area is now defined. If the feature is smaller than the local area and does not extend to the boundary then its extent is completely within the local area boundary and this extent defines the feature area size.

If the feature extends to the boundary of the local area and contains a portion in a neighbouring local area then these areas need to be grouped together into a larger feature area. The prerequisite for joining two local areas is that the feature area within each local area has a common border based on a four neighbour rule. The feature area extent is then extended to the limits given by the feature local areas in the feature area. The joining of feature local areas is done in a recursive manner, looking at the four neighbouring areas in turn until no neighbouring areas remain. This defines the feature areas.

The comparison of the two methods for the determination of the feature extent within local areas reveals that the line statistics method is unsuitable for the discriminating of separate features. The adaptive threshold method is ideally suited to this task and forms the basis for the feature discrimination process.

From the feature area a threshold is calculated to form a binary image of the feature area. The thresholds of the local areas that make up the feature area vary over a range and an optimal value is sought to threshold the feature area as a whole. This binary image often contains several "blobs" which may be features or may be growth rings or needle trace. Each blob is seperately identified and examined for candidature of a feature class. The results of this feature discrimination process are presented in Appendices 3A, 3B, 4A and 4B.

10.2.1 DETERMINING THE EXTENT OF FEATURES WITHIN FEATURE AREAS

A feature area may contain one or more features within its area. If the calculation of the extent of the feature within the local areas has resulted in a large feature area, for example the two features in WOOD01 (figure 10.13), then the extent of the feature within the larger feature area needs to be calculated. Because neither of the two methods of the first part of this chapter are able to reliably define the extent of the feature it is necessary to recalculate the feature extent.

A feature area can contain more than one feature as in WOOD42 (figure 10.14). As well as determining the extent of the feature within a feature areas the separation of the individual features must be considered. The task is to separate the individual features so that each can be classified into catagories.

The use of an adaptive threshold method, similar to that used for local areas, proves most effective in forming a binary image of the feature area, with individual features identified as separate blobs. A great deal of experiment with the parameter values of the peak-finding routine and the flow of analysis results in the flow diagram of figure 10.15. Several differences are present in the analysis of feature areas as opposed to local areas. These are discussed in relation to the flow diagram for the determination of the threshold value for feature areas.

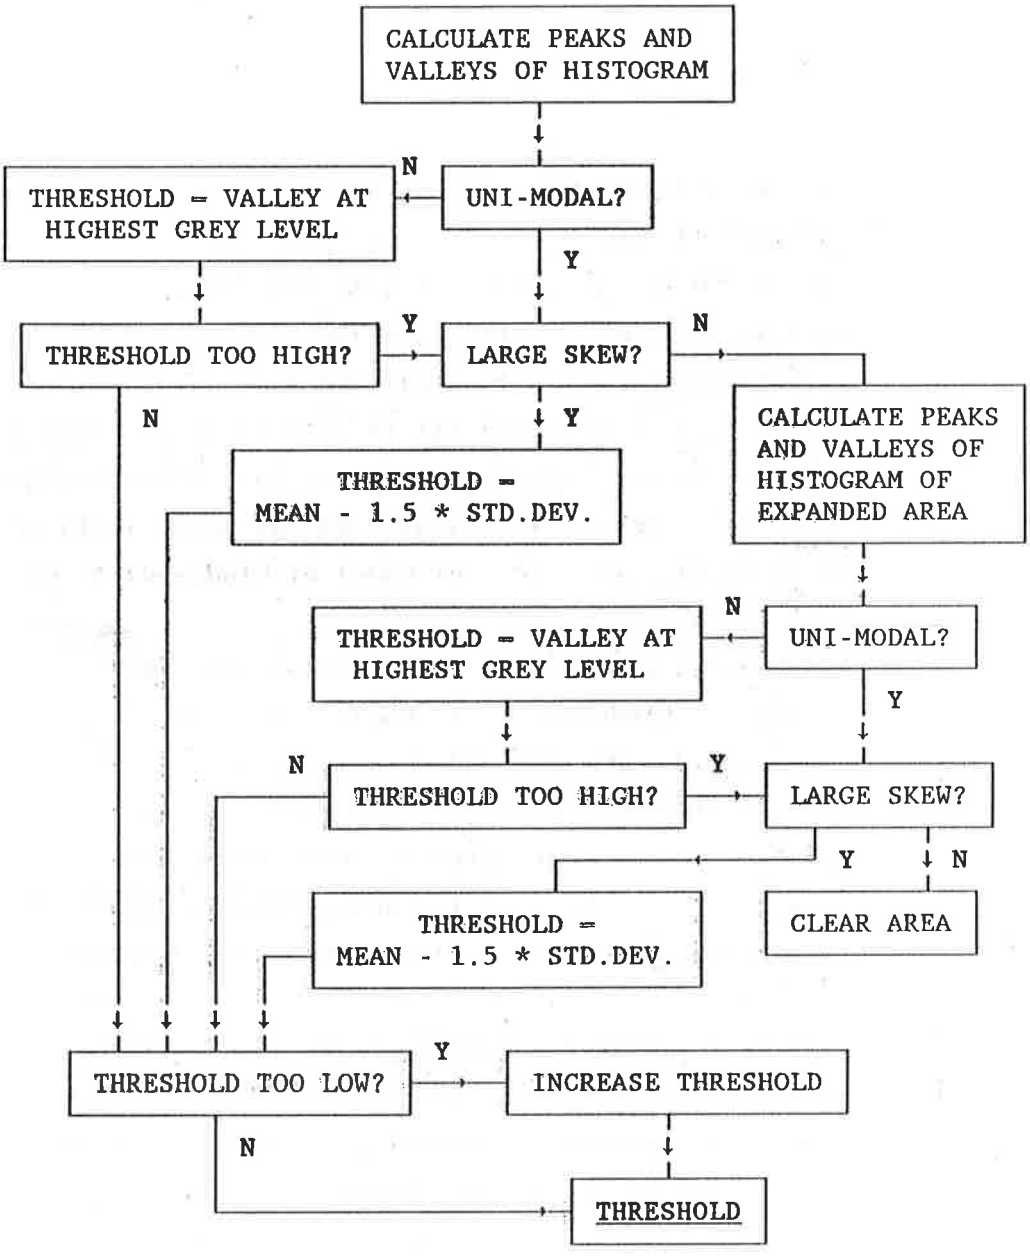

Figure 10.15 Flow diagram for determining the threshold value for variable size feature areas based on the grey level histogram.

The peak-finding routine is applied to the feature area with values for breadth and depth of 10 grey levels and 8 pixels respectively. These higher values suppress all but the largest peaks. The subsequent interpretation of the histogram is guided by several insights. The feature area can vary from a very small

area that totally defines a dark knot to a large area that contains a diffuse spike knot, dark growth rings, bark, and a great proportion of background area because the spike is at an angle.

If the area contains a significant proportion of background then the histogram is usually bi-modal and the threshold at the lowest point in the valley is optimal. If the histogram has more than one valley, the one at the highest grey level (below a suitable upper limit) is chosen because the area is most likely to have little background area. The upper limit is defined as 20 grey levels below the average of the means of all the clear areas in the image, referred to as the global clear mean. A threshold is found in this way in 26 of the 62 feature areas.

The absence of a suitable valley presses the requirement for an alternative method for threshold determination. If the feature area contains dark bark the histogram has a large skewness that is used to select a threshold in a similar manner to that for local areas. A value of

threshold = mean - 1.5 * s.d.

is found to give good results in the 24 feature areas that call for it.

In the examination of local areas, if a threshold cannot be found from a valley or the skewness, neighbouring local areas are inspected. The equivalent procedure in the treatment of feature areas is to expand the size of the feature area in all four directions (where this is possible). The feature area is increased by 16 pixels in each direction and the process of the search for a threshold is repeated. If no suitable threshold is found in this second search the area is labelled as clear.

The area is expanded in the examination of 16 of the 62 feature areas, where the histogram of the feature area is unimodal without a large negative skewness, and a threshold is found in 9 of these. The remaining 7 are classed as clear.

The lower limit for a threshold is set at 70 grey levels below the global clear mean. This is found to give a good

separation of the feature from the background in most cases without the inclusion of noise from dark growth rings. This increase of the threshold to a minimum is performed in 28 feature areas.

10.2.2 SEPARATING FEATURES WITHIN THE FEATURE AREA

Thresholding of the feature area creates a single bit plane mapping of the features in the feature area in the form of a number of blobs. A blob is simply a group of connected pixels. The feature area sometimes contains a single feature but often contains more than one feature. Most images after thresholding contain noise in the form of small disconnected blobs.

In order to separate the features in a feature area the blobs are identified individually by assigning each blob to a single separate grey level. Blob pixel connection is defined by the four neighbour connection rule, that is pixels of the same grey level and horizontally or vertically adjacent are neighbours and so part of the same blob. Diagonally adjacent pixels are not connected. A necessary consequence of this definition is that the background is defined by the eight neighbour connection rule. Background pixels are connected by all eight adjacent pixels that are of the same grey level.

When thresholding of the feature area reveals several blobs, each blob is catagorised separately. Blobs that contain less than a threshold number of pixels are classed as noise and removed. A minimum of 50 pixels describes a circular feature of 3.6 mm diameter or a rectangular feature 1 mm wide and 10 mm long. This is the minimum size of feature that is detected.

10.3 APPLICATION OF FEATURE DISCRIMINATION METHOD

The feature discrimination method has been described as it has been developed with application to the training set of the WOOD library of images. It remains to assess the effectiveness of the method for the test set of images as well as the training set. The entire procedure is as follows:

Each image in both the training and test sets are divided into local areas of size 64 * 64 pixels. A decision boundary is derived using the training set of images using the four tonal measures of mean, standard deviation, skewness and kurtosis.

FEATURE DETECTION

For each local area in each image the tonal measures are calculated and the decision boundary is applied which assigns a class of feature or clear to each local area.

FEATURE DISCRIMINATION

Line statistics method - For each local area the line mean and line variance is calculated for each row/column, starting from the edges and working inward, until the either the line mean falls below the mean threshold or the line variance falls above the variance threshold. This marks the extent of the feature within the local area.

Adaptive threshold method - For each feature local area a threshold is determined using the information contained in the grey level histogram. This threshold divides the local area into feature and background and allows the extent of the feature within the local area to be estimated.

If the feature extends across adjacent local areas they are joined together to form feature areas. A threshold is determined for each feature area in a manner similar to the adaptive threshold method. The threshold produces a number of feature blobs which give the location and shape of the features within the feature area.

The results of the application of feature detection and feature discrimination is presented in the figures of Appendix 3A and 3B, for the line statistics method, and Appendix 4A and 4B, for the adaptive threshold method. The original images appear in Appendix 1. The decision boundary that classifies the local areas

uses the four tonal measures and results from 2000 iterations of the classification algorithm applied to the training set:

$$w = (135.7, -146.1, 101.3, -138.7, 5.84)$$

The feature discrimination process is that described in this chapter and illustrated in figure 10.8.

Each of the Appendices 3A, 3B, 4A and 4B contain a list of each blob in each feature area and a number of shape measures to describe the blobs. The blobs that correspond to features are labelled to enable the reader to associate the blob image with the shape measures.

The last part of the four appendices contains the features listed in the feature catagories of table 9.2 with a description of how well the feature is isolated from the background. For each feature is listed the size, DY', that is the measured size of the blob. The size, DY, was measured by an experienced Quality Control Officer at the time that the image was recorded (as described in chapter 5).

10.3.1 MEASURES FOR THE CLASSIFICATION OF FEATURE BLOBS

Examination of Appendix 3A, 3B, 4A and 4B reveals that there are a lot of blobs that are associated with dark growth rings, bits of bark, etc. The final task of the discrimination process is to classify the blobs into the catagories of features required as input to a grading process. This is initially investigated by examining the measures of each blob to determine if any single measure can be used to specify uniquely one or more feature types.

Several important conclusions can be drawn from the feature lists in the appendices. Consider table 10.2a and 10.2b which lists the sizes of the known features from the size of the blobs.

SIZE	CORRECT	OVER	UNDER	LOST	TOTAL
LARGE > 30mm	21	-	7	1	29
MEDIUM 20-30mm	13	8	1	2	24
SMALL 10-20mm	27	5	5	2	39
PIN <10mm	19	3	-	2	24

Table 10.2a Numbers of knot features that are given sizes by the feature discriminator. Line statistics method is used to calculate local area extent.

SIZE	CORRECT	OVER	UNDER	LOST	TOTAL
LARGE > 30mm	23	-	4	2	29
MEDIUM 20-30mm	10	7	2	5	24
SMALL 10-20mm	27	5	4	3	39
PIN <10mm	18	4	-	2	24

Table 10.2b Numbers of knot features that are given sizes by the feature discriminator. Adaptive threshold method is used to calculate local area extent.

Of the 116 knot features present (in both the training and test sets) the reasons for a feature being undefined are:

a) - the feature local area is undetected by the decision boundary (a medium knot in WOOD13, a small knot in WOOD20, and a pin knot in WOOD14).

b) - the local area extent is unable to be defined. This only occurs with the adaptive threshold method on one occasion with a medium knot (WOOD15).

c) - the feature area threshold is unable to define the feature. This results in the entire feature area being classed as clear 2 times with the line statistics method and 4 times with the adaptive threshold method (plus a fifth that is a false positive area).

Medium size knots are most likely to be lost. An inspection of the original knot images reveals these features to be faint, intergrown knots that do not possess much contrast to the background. It is not surprising that these features are not easily defined by a single threshold value. In addition to the features that are lost, a further 12 are defined only by the bark that partially encases the knot, the knot being undefined. This only gives a good indication of the size of the knot in 2 cases. All these latter 12 are either large or medium knots.

Small knots and pin knots are for the most part well defined. This is because small and pin knots are usually of high contrast to the background and they also interact with growth rings far less than the larger knots. Medium and large knots exhibit a much larger variation between the measured size and the blob size primarily due to the effect of growth rings extending the blob (7 instances).

The size, as measured by DY', is a shape factor for blobs. It is a crude measure, however, and not sufficient to distiguish between different types of features. Another simple measure is the length/width which is used to distinguish between round features (knots) and long thin features (pith, resin and bark). The figures in table 10.3 reveal that this distinction is quite clear and separate. All knots and holes have a length/width ratio of less than 2.8 and all resin, bark and pith features have a ratio greater than 3.7. (One pin knot that includes a growth ring that is excluded from the table has a ratio of 3.54 and so does not violate the separate nature of the measure.)

FEATURE TYPE	RANGE	
	LINE STATS. METHOD	ADAPT. THRESH. METHOD
KNOTS AND HOLES	0.47 - 3.04	0.39 - 2.79
PITH	3.87 - 14.11	3.87 - 14.11
RESIN AND BARK	3.38 - 22.11	3.75 - 22.11

Table 10.3 Ranges of the measure **LENGTH/WIDTH**. One pin knot in WOOD08 is ignored because it includes a growth ring.

A non-dimensional shape measure is perimeter^2/area and the ranges for the different features are listed in table 10.4. The lowest value for this measure is given by a circle:

$$\frac{(\text{PERIMETER})^2}{\text{AREA}} = \frac{\Pi^2 D^2}{\Pi D^2/4} = 12.6$$

FEATURE TYPE	RANGE	
	LINE STATS. METHOD	ADAPT. THRESH. METHOD
LARGE KNOTS AND HOLES	36 - 464	39 - 455
MEDIUM KNOTS AND HOLES	22 - 180	19 - 200
SMALL KNOTS AND HOLES	19 - 115	19 - 79
PIN KNOTS AND HOLES	20 - 36	21 - 39
PITH	80 - 130	81 - 130
RESIN AND BARK	41 - 345	41 - 385

Table 10.4 Ranges of the measure **PERIMETER^2/AREA**. One pin knot in WOOD08 is ignored because it includes a growth ring.

The measure perimeter^2/area is not sufficient to distinguish features into classes but does provide useful information about the blob. A large value for the measure indicates that the blob is extended in some way. Pith can naturally be expected to have a high value because it is long and thin as is resin and bark. Knots, particularly medium and large knots, have large values because of the presence of growth rings which increase the value of the perimeter. Small knots and pin knots are less affected by growth ring extensions of the blob perimeter and the figures in table 10.4 indicate that pin knots can be distinguished from pith, resin and bark while small knots can be distinguished only from pith using this measure alone.

10.3.2 THE INFLUENCE OF TEXTURE IN THE FEATURE DETECTION STAGE ON THE RESULTS OF FEATURE DISCRIMINATION

In the previous chapter it is found that texture measures can significantly improve the classification of local areas. It is now possible to determine the effect of this improvement on the discrimination of features.

Table 9.18 lists 29 local areas of the training set that are classified alternatively when one or another of five texture measures is added to the tonal measures. These measures are contrast (CON), correlation (COR), local homogeniety (LHOM), cluster shade (CSH) and cluster prominence (CPR). CSH and CPR are correlated to the extent that they classify all the local areas identically and so are considered together.

Often the correct classification of a local area, for example area (3,4) of WOOD01, does not lead to a better definition of the feature. In this case the area extends the feature area to encompass the whole of the knot but the feature discrimination process still thresholds the area such that only the bark is isolated. The discrimination process that relies on the form of the grey level histogram is unable to make use of the textural information.

Similarly, COR classifies area (6,4) of WOOD01 as a feature area when the given classification is clear. Close

inspection of the local area reveals the presence of a very faint feature, so faint that it is given a clear classification by the author in the compilation of the local area classifications. Although COR detects the presence of a feature the feature discrimination process is unable to make use of the information and the area is classed as clear.

Many of the local areas end up having no or little effect on the discrimination of the feature and so only those that have a significant effect are discussed. Local area (5,1) of WOOD09 is correctly classed as a clear area with all the texture measures under consideration. The result of this is to break up a section of a resin/bark pocket into 2 separate areas. The significance of correctly classifying area (6,2) is even more significant as this area connects the two features in the image into a single feature area. By considering the 2 features in separate areas the discimination process is able to reach a better threshold value for each area. The result is that the features are better defined.

WOOD10 contains several areas that contain strong vertical growth rings associated with a large knot.Areas (4,5) and (5,2) are correctly classed as clear with the addition of any one of the texture measures and (5,2) is particularly significant as it ends up being a large growth ring blob after the feature discrimination process. Correct identification of this area removes a large noise blob. Area (3,4) contains a dark growth ring that is connected to a large knot. The incorrect classification of this area extends the feature area slightly. The incorrect classification of area (5,4) is more significant as it results in the connection of the two main feature areas in the image into one large feature area. The threshold determined by feature discrimination is not as good as that obtained for each of the two areas. CSH and CPR do not make this misclassification.

The large spike knot of WOOD21 is defined by two feature areas. Correctly classifying either area (1,3) using CON, COR, or LHOM, or area (2,2) using CSH results in the two areas being

connected into one feature area. The threshold of this area leads to a better definition of the feature.

This investigation of the significant improvement afforded by using texture information in the feature detection stage shows that cluster shade and cluster prominence give the best results. The ability of the texture information to aid in the detection of faint knots is significant. However, the feature discrimination process is based on the statistics of the grey level histogram of the feature area and so cannot make use of this texture information. Texture information is required in the discrimination stage to be able to better define the features.

10.3.3 CONCLUSIONS OF FEATURE DISCRIMINATION

The use of the grey level histogram to define the size of features, either line by line with the line statistics method or over the whole local area with the adaptive threshold method, is satisfactory for features that contrast with the background. This includes pith, resin, bark, and most pin and small knots. Medium and large knots are more likely to have a strong pattern of growth rings that surround the feature in such a way as to make a size definition impossible using a grey level threshold. This difficulty is not restricted to an image analysis program. Some of these features are difficult for human inspectors to size without ambiguity and the aspects that are used are textural in nature. In particular, the extent of an intergrown spike knot is taken as the largest complete knot growth ring. The next growth ring is considered part of the background.

Textural and contextual information is needed to discriminate between the edge of medium and large knots and the growth rings that surround them. Tonal measures will never be able to define these diffuse features adequately. Tonal measures are adequate for the discrimination of pith, resin and bark and can be used to identify these features as distinct from knots using simple blob shape measures. This restricts the need for textural measures to a smaller, more specific set of features, reducing computational load.

11. GRADING TIMBER PLANKS

The output of the feature discrimination process is a list of the features in an image of radiata pine. The list also contains the size and location of each feature detected. Assuming that this list is an accurate representation of the actual plank then it is a straighforward process to apply the grading rules to place the plank into a grade defined by the Australian Standards.

The rules that apply to sawn boards are listed in Appendix 5. They are codified into a program that generates a symbolic image of the plank from a list of features. Seven examples are shown in the following figures along with the grading program output to illustrate the range of feature distribution among the different grades.

Each plank is 5.99 metres long and is represented in the figures by one metre lengths. Encased knots are red with a black circle around them. Only one face is graded.

The important point to be stressed is that the transfer of the grading rules into a computer program is a simple process. The program written for this chapter reads from a list that contains the location, size and type of every feature on a plank. This is the output of the feature discrimination process. The program performs the following steps:

1 - Check if docking is required.
2 - Check suitability for standard grade.
3 - Check suitability for select grade.
4 - Check suitability for joinery grade.
5 - Check suitability for clear grade.

The program stops when it is found to have features that exclude it from a grade.

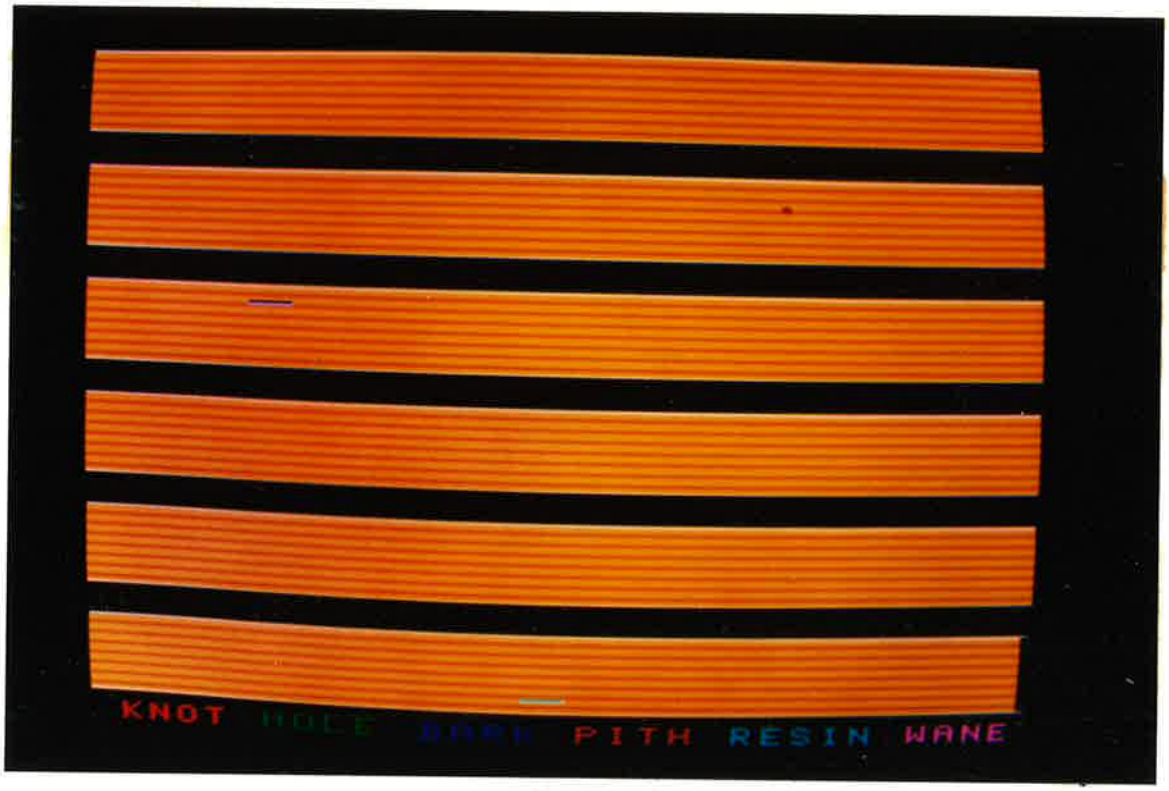

Figure 11.1 Top end of joinery grade.

```
FEATURE LIST FILE IS board1
THERE ARE    3 FEATURES
DOCKING NOT REQUIRED
SUITABLE FOR STANDARD GRADE
SUITABLE FOR SELECT GRADE
SUITABLE FOR JOINERY GRADE
BOARD IS NOT SUITABLE FOR CLEAR GRADE
```

Clear grade timber is free from all features and is used in the production of mouldings which are special small sections such as picture frames. The presence of any feature will exclude the piece from clear grade and move it into the next highest grade, joinery. Joinery timber is used for items such as door and window frames which require a straight timber with few knots and strict limits apply to holes.

Figure 11.2 Bottom end of joinery grade.

```
FEATURE LIST FILE IS board2
THERE ARE  18 FEATURES
DOCKING NOT REQUIRED
SUITABLE FOR STANDARD GRADE
SUITABLE FOR SELECT GRADE
SUITABLE FOR JOINERY GRADE
BOARD IS NOT SUITABLE FOR CLEAR GRADE
```

This piece is on the limit of what is acceptable for joinery grade. The three holes are within size limits but would exclude the piece from joinery grade if they are any closer together. Similarly the size and distribution of the knots, bark and resin place it on the lower end of what is acceptable.

Figure 11.3 Top end of select grade.

```
FEATURE LIST FILE IS board3
THERE ARE  23 FEATURES
DOCKING NOT REQUIRED
SUITABLE FOR STANDARD GRADE
SUITABLE FOR SELECT GRADE
    1 LENGTH OF PITH - UNACCEPTABLE
   ENCASED KNOT TOO LARGE - 2 OCCURRENCES
   HOLE TOO LARGE :   11
   MORE THAN ONE HOLE IN TWO METRES - 3 OCCURRENCES
   KNOT TOO LARGE FOR JOINERY GRADE
    8 OVERSIZE RESIN OR BARK - UNACCEPTABLE
   BOW EXCEEDS LIMITS FOR JOINERY TIMBER - 168
   SPRING EXCEEDS LIMITS FOR JOINERY TIMBER -  67
BOARD NOT SUITABLE FOR JOINERY GRADE - 18 DEFECTS
```

The presence of any one of the defects listed exclude the piece from joinery grade. Select grade is used where appearance and strength are both important. Uses include panelling and shelving where a knotty timber with a fine appearance is sought.

Figure 11.4 Bottom end of select grade.

```
FEATURE LIST FILE IS board4
THERE ARE  34 FEATURES
DOCKING NOT REQUIRED
SUITABLE FOR STANDARD GRADE
SUITABLE FOR SELECT GRADE
    1 LENGTHS OF PITH - UNACCEPTABLE
   ENCASED KNOT TOO LARGE - 2 OCCURRENCES
   HOLE TOO LARGE :   11
   MORE THAN ONE HOLE IN TWO METRES - 3 OCCURRENCES
   KNOT TOO LARGE FOR JOINERY GRADE - 4 OCCURRENCES
   MORE THAN THREE KNOTS IN FOUR METRES
    6 OVERSIZE RESIN OR BARK - UNACCEPTABLE
   BOW EXCEEDS LIMITS FOR JOINERY TIMBER - 168
   SPRING EXCEEDS LIMITS FOR JOINERY TIMBER -  67
BOARD NOT SUITABLE FOR JOINERY GRADE - 20 DEFECTS
```

Select grade will accept a considerable quantity of features as long as they fall within the requirements of size and distribution.

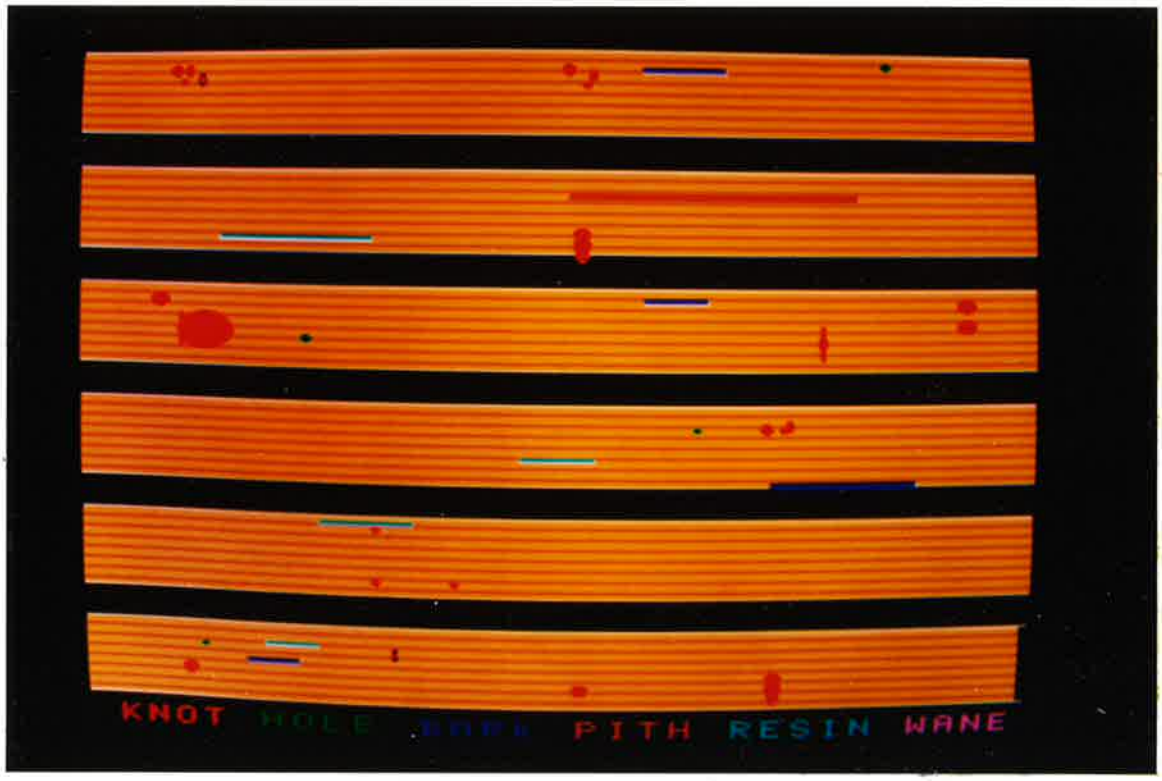

Figure 11.5 Top end of standard grade.

```
FEATURE LIST FILE IS board5
THERE ARE  42 FEATURES
DOCKING NOT REQUIRED
SUITABLE FOR STANDARD GRADE
    5 OVERSIZE RESIN OR BARK - UNACCEPTABLE
BOARD NOT SUITABLE FOR SELECT GRADE -  5 DEFECTS
```

Timber in this grade has many of the characteristics of select and even joinery grade but is less restrictive. It is suitable where the requirements for finishing are less exacting such as flooring. Many pieces are finished on one side and have more numerous or larger features on the reverse side.

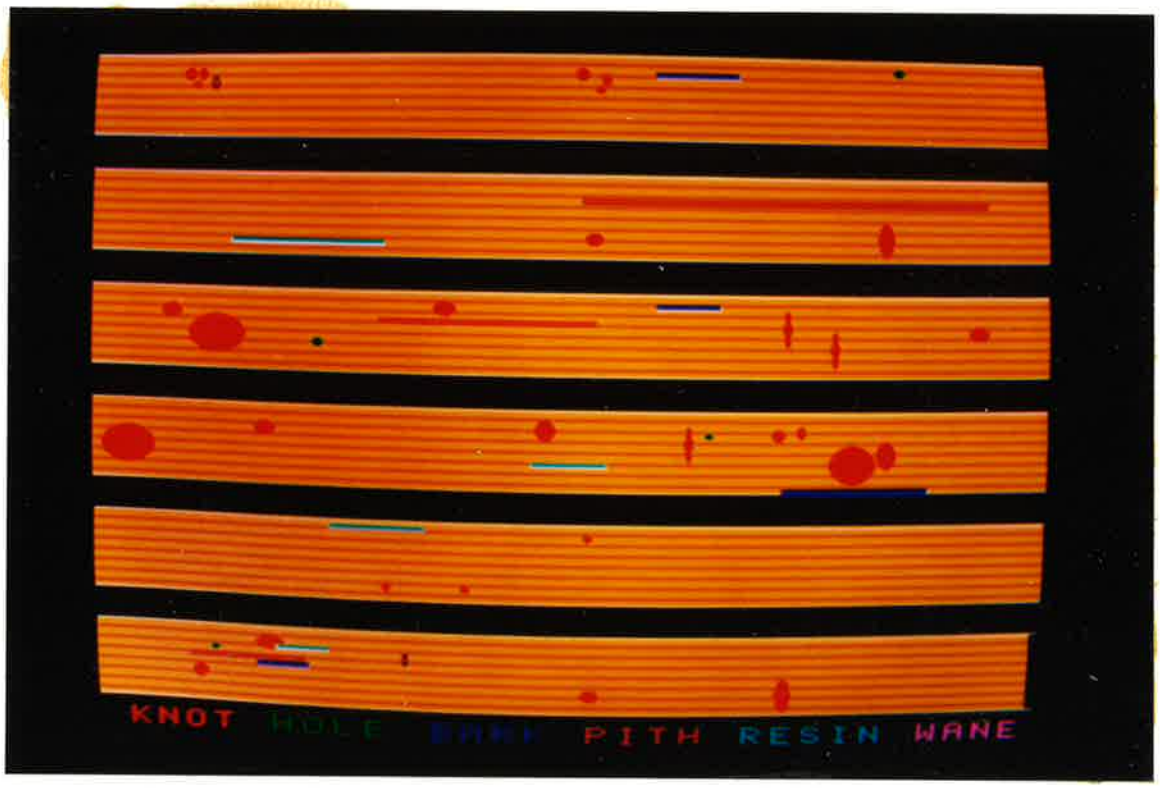

Figure 11.6 Bottom end of standard grade.

```
FEATURE LIST FILE IS board6
THERE ARE  49 FEATURES
DOCKING NOT REQUIRED
SUITABLE FOR STANDARD GRADE
   763.5 PITH TOO LONG - UNACCEPTABLE EXCEPT INFREQUENTLY
    5 OVERSIZE RESIN OR BARK - UNACCEPTABLE
BOARD NOT SUITABLE FOR SELECT GRADE -  6 DEFECTS
```

This represents the limit for standard grade material with seveal knots approaching the maximum permissible size of half the face width and the pith approaching the maximum allowable length.

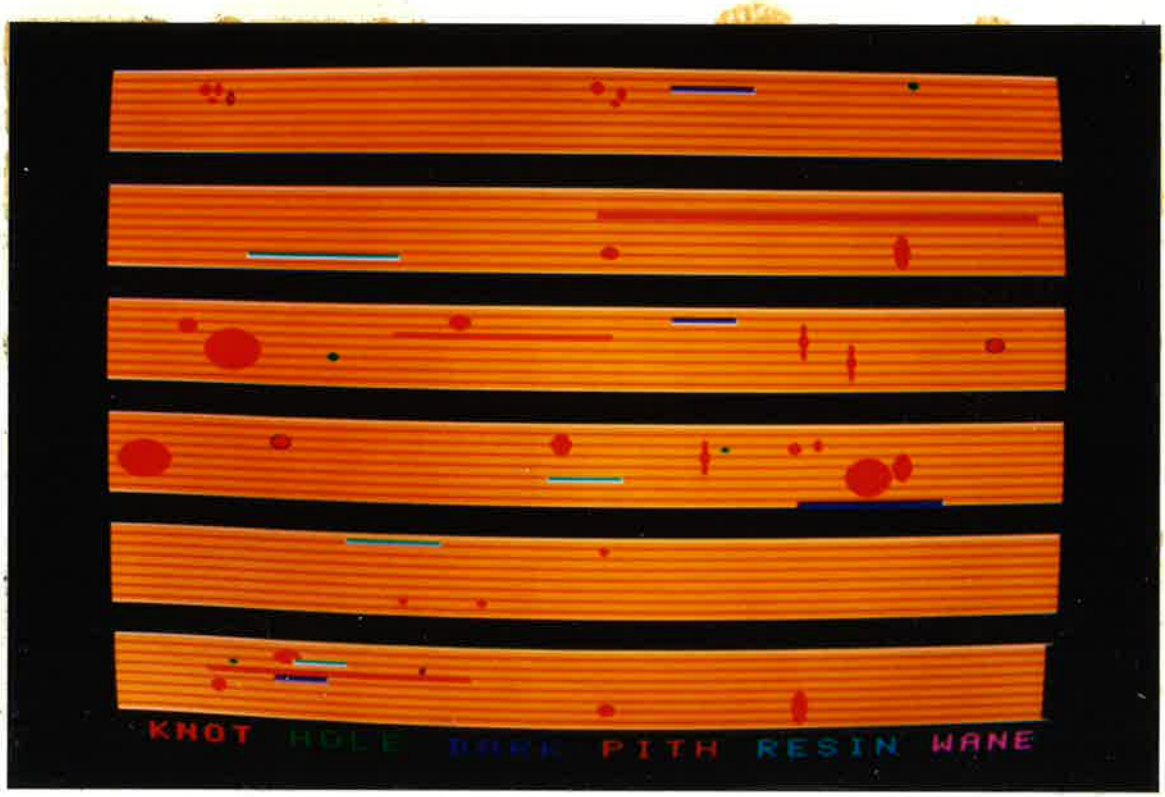

Figure 11.7 Merchantable grade.

```
FEATURE LIST FILE IS board7
THERE ARE  49 FEATURES
   DOCKING REQUIRED AT  2096 AND  2186  KNOT IS   62 WIDE
BOARD REQUIRES DOCKING
   464.5 INDIVIDUAL PITH TOO LONG - UNACCEPTABLE
   953.5 PITH TOO LONG - UNACCEPTABLE
   ENCASED KNOT TOO LARGE - 2 OCCURRANCES
   KNOT TOO LARGE FOR STANDARD GRADE
BOARD NOT SUITABLE FOR STANDARD GRADE -  5 DEFECTS
   953.5 PITH TOO LONG - UNACCEPTABLE EXCEPT INFREQUENTLY
   ENCASED KNOT TOO LARGE - 2 OCCURRANCES
   KNOT TOO LARGE FOR SELECT GRADE
    5 OVERSIZE RESIN OR BARK - UNACCEPTABLE
BOARD NOT SUITABLE FOR SELECT GRADE -  9 DEFECTS
BOARD SUITABLE FOR MERCHANTABLE GRADE
```

Any one of the defects listed will exclude the piece from standard or select grade. Merchantable grade timber is that which falls outside all the grades specified in the Standards but can still be sold. Note the differences in the lists of defects between select and standard grades.

12. DESIGN OF A PRODUCTION SYSTEM

The purpose of this study has been to demonstrate the feasibility of methods that can be used to automate the visual inspection of timber. This is done using an off-line environment in order to test several different ideas. No attempt is made to perform the inspection task at production speeds although the development of a real-time system is expected to result from this work. With this longer term goal in mind the methods that have been presented in this study are amenable to incorporation into a parallel architecture where the required speed of processing can be achieved by multiplying the number of processors.

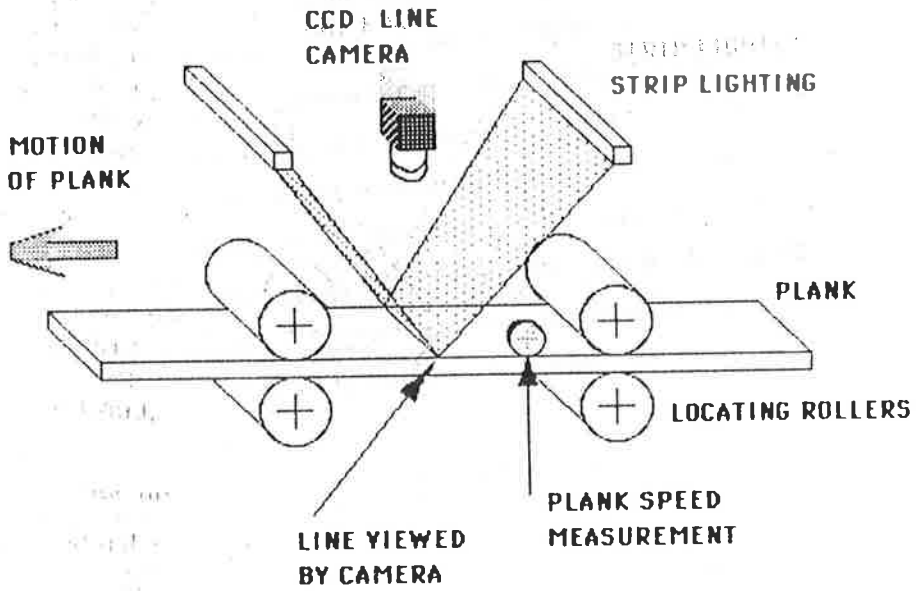

Figure 12.1 Setup of an automated visual inspection system.

This section outlines the way in which the methods of this study can be developed into a real-time production system for the purpose of automating the visual inspection of timber.

12.1 IMAGING HARDWARE

The camera used in this study is an array camera that images a rectangular section of wood onto a rectangular array of sensing elements. If this were to be used in a production system there would exist the problem of aligning consecutive frames of the plank into a continuous representation. This can be overcome by using a linear array camera that consists of a single row of sensing elements. A roller that is in contact with the wood can trigger the camera to capture a row of pixels at equal distances along the length of the board (figure 12.1). These would be stored in an enlarged memory that can be conceived of as an array representing the whole plank (figure 12.2).

Images are processed in this study a single frame at a time and require repeated transfers of data from the image memory to the system memory and back again. This is the result of the interactive nature of the development process. A production system will process the image data continuously by using dual ported static random access memory (SRAM) to enable the camera to transfer data into memory at the same time as the analysis processors are reading the image data.

The plank speed measurement device consists of a roller with an optical encoder that detects the presence of the plank and its speed. This information is used by the controller to trigger the linear camera at equal linear distances along the plank. A data switch places the camera data into consecutive locations in memory.

The feature detection processor reads the image data and processes it in local area size arrays that correspond to rectangles of the plank to calculate the appropriate feature measures. The decision boundary is applied to the feature measures of each local area and if it is classified as a feature area the location address is placed in the FIFO (first in first out). This decouples the feature detection processor from the feature discrimination processor. The controller ensures that the

feature detection processor does not try to pass the position in memory to which the camera is writing.

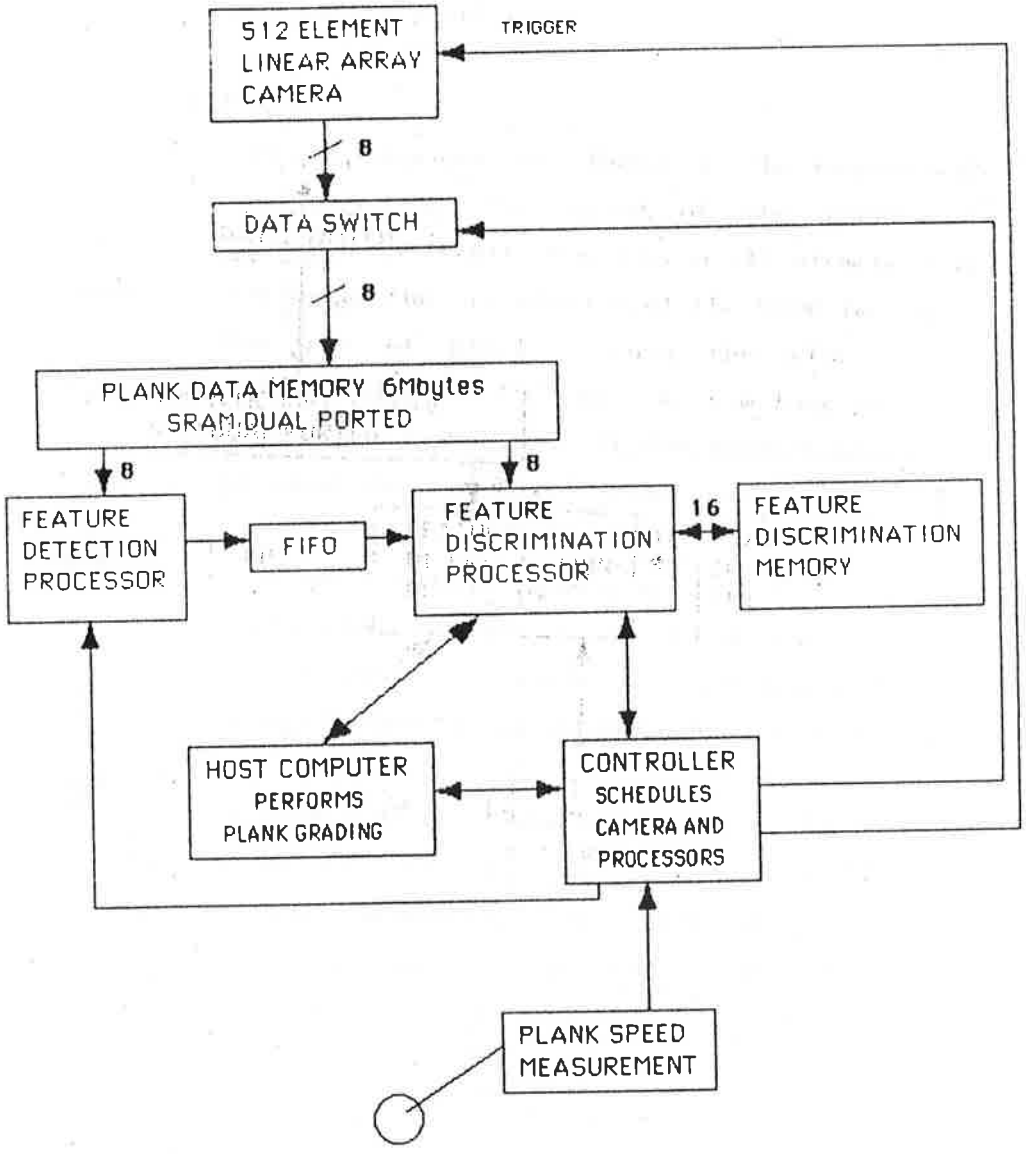

Figure 12.2 Schematic diagram of a possible real-time automated visual inspection system for radiata pine.

The feature discrimination processor will read the location of feature areas from the FIFO and read the appropriate image information into its own memory. It will then perform the computations necessary to merge feature areas together and will

form a list of attributes of each feature area and make a decision on the type of feature that has been detected. A list of the features together with their locations is passed to the host processor which places the plank into the proper grade when the information for the entire plank has been collected.

12.2 PROCESSING SPEED

To fulfill the demands of industry the inspection is required to be performed at the speed of the proposed new finishing machines which is 3 m/s. Consider a 512 element linear array camera operating at the resolution of the WOOD images used in this study. One row of pixels across the plank at 0.376 mm/pixel gives a plank width of 192 mm. A 6 m long plank at a resolution of 0.533 mm/pixel requires 11,260 rows of pixels. The total memory requirement for an entire plank (where one pixel is one byte or 256 grey levels) is 5.8 Mbytes.

The rate at which this data must be stored in memory and processed is determined by the plank speed of 3 m/s. The data rate is therefore 2.9 Mbytes/s which is well below the video digitising rate of the MATROX frame-grabbing board which is more than 6.5 Mbytes/s.

64 rows are stored from the camera every 11.3 ms and make up 8 local areas of size 64 * 64 pixels. Each local area requires 4096 integer additions, to fill the grey level histogram, 256 floating point operations (FLOPS) for the normalisation, 256 FLOPS for the mean, 512 FLOPS for the standard deviation, and another 512 FLOPS to calculate the skewness and kurtosis. This is a total of 12,288 FLOPS in 11.3 ms or 1.1 MFLOPS/s. This is well within current technology for a single processor.

The feature discrimination processor does not have to process the entire area of the board but does have to perform more varied computations than the feature detection processor. This processor operates on one feature area at a time, formed from connected feature local areas, and the speed of this may be limiting. However, as each feature area is separate from all

others the analysis of these areas can be distributed to a number of processors by the controller. This parallelism can be extended as far as is necessary to process the maximum expected number of features per board. This is dependent upon the speed of processing to discriminate a feature, and the number of features expected in a plank. The former is dependent upon the algorithms used and the hardware that they are implemented upon. The latter can be estimated from tables 12.1 and 12.2.

SIZE	SOUND KNOTS	ENCASED KNOTS	HOLES	SPIKE KNOTS
LARGE (> 30mm)	137 (0.792)	53 (0.306)	7 (0.040)	134 (0.775)
MEDIUM (20-30mm)	431 (2.490)	109 (0.630)	21 (0.121)	8 (0.462)
SMALL (10-20mm)	735 (4.250)	234 (1.350)	55 (0.318)	-
PIN (< 10mm)	1258 (7.270)	299 (1.730)	63 (0.364)	-
PITH	190 lengths (1.098), average length = 290mm			
SPLITS	11 lengths (0.064), average length = 181mm			
WANE	8 lengths (0.046), average length = 429mm			

Table 12.1 Number of features in a production run of 190*19 appearance grade timber. Total length = 1038 m. Figures in brackets are the probabilities of occurrence in a 6 m length. Total area = 197.2 m^2.

SIZE	SOUND KNOTS	ENCASED KNOTS	HOLES	SPIKE KNOTS
LARGE (> 30mm)	41 (0.158)	2 (0.008)	3 (0.012)	97 (0.373)
MEDIUM (20-30mm)	181 (0.696)	34 (0.131)	10 (0.038)	14 (0.054)
SMALL (10-20mm)	404 (1.550)	82 (0.315)	36 (0.138)	8 (0.031)
PIN (< 10mm)	564 (2.170)	160 (0.615)	73 (0.218)	-
PITH	31 lengths (0.119), average length = 249mm			
SPLITS	49 lengths (0.188), average length = 200mm			
WANE	negligible			

Table 12.2 Number of features in a production run of 42*19 appearance grade timber. Total length = 1560 m. Figures in brackets are the probabilities of occurrence in a 6 m length. Total area = 65.5 m^2.

Tables 12.1 and 12.2 were compiled from video tape of the mill output which was recorded at the Mount Gambier mill of the South Australian Woods and Forests Department. Several aspects of these tables are worthy of note. Although the timber in table 12.1 contains 3 times the surface area of the timber in table 12.2 it contains somewhat less than 3 times the number of features. One exception is large encased knots, which are of the order of size of the smaller section and lead to the separation of the length either by breaking or by docking. Another is pith which is more frequent in the wider planks. This is to be expected because the wider sizes can only be cut from the centre portions of the log and so are more likely to show pith on the

face. Wane also occurs more often on the wider size with none found on the small size for the sample inspected.

Knowledge of the frequency of occurrence of various features can be used to optimise the feature discrimination software by directing the search for probable feature classes in the most likely areas.

12.3 OTHER IMAGING SYSTEMS

The results of this study indicate that direct overhead lighting is unable to discriminate features that contain depth such as splits, cracks, knot chipping, skip and some holes. Holes can be detected by backlighting, that is, by illuminating the plank from underneath. If a dual head system is employed to grade both faces of the plank at the same time then the back light will merely be the direct light from the opposite head. A separate linear array camera will be used to detect holes by having the gain of the detector reduced and a threshold element implemented. The information would be binary and transfer directly to the host processor to modify the feature list.

Side lighting is used enhance height information by producing shadows. The wood is illuminated by a source at a shallow incidence from a direction perpendicular to the direction of motion. Splits and cracks are usually parallel to the plank length and will appear as a discontinuity to a third linear array camera, this time set with a high gain to reduce confusion from elements of dark wood. A separate processor would compile this information and pass it to the host processor to be included in the list of plank features.

12.4 PLANK LOCATING SYSTEM AND LIGHTING

Operation of the visual inspection system relies on the correct positioning of the timber with respect to the cameras. Unobstructed access is required to the plank from above and below. The plank must be positioned in the focal plane of the cameras and computation is simplified if a consistent datum edge can be ensured.

Two pairs of rollers at the entry to the visual inspection area ensure that the plank enters in a horizontal attitude. Side rollers with a degree of compliance force the plank laterally against the locating datum edge which ensures that each camera scan starts on the plank.

Direct lighting in the form of strip lights, as in figure 12.1, will ensure an even illumination of the wood surface. The use of two strip lights will provide a degree of backup should one light fail during service. If the camera uses an automatic gain control production can continue while the failed unit is replaced.

12.5 PLANK DEFORMATION MEASUREMENT

An important aspect of grading radiata pine that is not a visual inspection is the measurement of plank deformation. Diagrams of the plank deformations of bow, spring, cupping and twist are shown in figure 2.8 and are due to differential contraction during the drying process.

Ideally a sensor is required that is non-contacting, so as not to mark the timber, quick and reliable with a resolution of the order of 1 millimetre. A sensor fulfilling these requirements is the ultra-sonic ranger. Using a transducer of 10 mm diameter and a frequency around 180 KHz produces a beam that disperses to around 20 mm diameter at a range of 150 mm. The attenuation of the sound wave is very large at this frequency and so the sensors can be operated without interference from each other. Background noise at this frequency is effectively zero although measurements would need to be performed to confirm this in the saw mill environment.

A single micro-processor can interrogate up to 200 of these sensors per second in a serial fashion. A possible design could have 7 sensors spaced at 50 mm spacing across the board at one end and repeated every 300 mm from 2.1 m up to 7.2 m. This is a total of 126 sensors to measure cupping, bow and twist. A further 18 sensors to measure the edge of the board will give the spring for a total of 144 sensors. This would require the board

to be positioned along a datum edge in a stationary position for less than one second. With a little ingenuity the number of sensors could be reduced and the measurements taken of the moving board at the exit of the visual inspection unit.

12.6 CONCLUSIONS

The construction of an automated system for the inspection of radiata pine is possible and would consist of the following features:

a) 3 linear array CCD cameras.
b) 6 Mbytes of dual ported static random access memory to store the image of an entire plank.
c) A single processor to perform feature detection.
d) A number of processors to perform feature discrimination.
e) Back lighting to enhance through the board features.
f) Side lighting to enhance cracks, splits and chipping.
g) Ultra-sonic range finders to measure plank deformation.

The total cost of this hardware is of the order of $35,000 and a system that could be implemented in a production environment would require a development time of approximately two to three years. A development program would require the services of the following team of skilled personnel:

a) A computer programmer/specialist to further develop the image processing routines and to encode them in assembler. Would also need to incorporate inputs from a number of other sensors into the controlling program.
b) An electronics technician to assemble the computer hardware devices into a reliable package.
c) A mechanical engineer to design and build the optics, lighting, plank deformation sensors, and plank locating system.

With these resources it is the authors opinion that a real-time automated visual inspection system for the grading of radiata pine boards can be developed.

13. THE FEASIBILITY OF AUTOMATING THE VISUAL INSPECTION OF RADIATA PINE

The overall conclusion that is evident from this thesis is that the automated visual inspection of radiata pine is a realisable task. The features can be detected at the production speed of a mill with a single processor. The discrimination of features can then be shared by a number of processors working in parallel to bring this stage up to production speed. The application of the grading rules is then easily performed.

The literature survey of chapter 3 reveals that automating the visual inspection of timber to meet a standard specification has not been attempted before the present work.

Various aspects of the camera and lighting are investigated in chapter 5 which reveals that there is no manual control of the camera gain. This forces the search for analysis routines for feature detection and discrimination that are isolated from overall brightness levels. This is considered to be positive from the point of view of designing a system that is robust in an industrial environment.

Chapter 6 outlines a preliminary investigation of the ability of tonal measures to describe the features on boards. The conclusion is that no single measure is enough to adequately describe the feature but that the concept of using local areas is useful in limiting the area of board to be processed.

Pattern classification is introduced in chapter 7 as a means of combining a number of feature measures in order to detect feature local areas. With this powerful tool texture measures of the local areas are added to the tonal measures in chapter 8 and an initial assesment of the performance of the classification algorithm is performed. This provides the groundwork for a more thorough investigation of the

classification algorithm and the effect of tonal and texture measures.

A detailed investigation of the classification algorithm parameters is carried out in chapter 9 and the values necessary for a stable iteration are determined. By calculating the ability of different combinations of tonal measures to detect the feature local areas the importance of each measure is assessed. Mean, standard deviation and kurtosis are all found to be essential for an adequate performance. Skewness, while not vital, does contain information that can be used at a later stage.

Texture measures are then added singly to a pattern space consisting of the four tonal measures to determine if the texture measure contains useful information that aids the detection of feature local areas. It is found that several texture measures improve the performance of the classification but only if the 0° direction of the SGLDM is used. This significant finding reduces to one quarter the computational burden associated with the calculation of the texture measure.

Two methods of determining the extent of features within local areas is investigated in chapter 10 using grey level statistics. Both methods prove capable of determining the feature extent for the purpose of joining local areas into feature areas with the adaptive threshold method providing the basis of the method for discriminating the features within a feature area. Shape measures are capable of distinguishing between feature types.

The shortcomings of the discrimination process are addressed in 10.3.2 where the influence of texture in the discrimination process is discussed. It is concluded that although tonal measures are adequate to detect and discriminate most features that texture measures are required, at least in the discrimination stage, to define the features sufficiently well for a production system to operate reliably. This is the direction that further research will take in the development of an industrial inspection machine.

The application of the grading rules is presented in chapter 11 with a visualisation of the distribution of features allowed by the different grades. The conclusion of this exercise is that the codification of the grading rules is a simple matter.

Chapter 12 considers the elements of the design of a production system. A considerable amount of development work is required to reach a fully automated system but this thesis demonstrates that such a system is capable of being realised with current technology.

It is the hope of the author that this system is developed and implemented in the near future.

APPENDIX 1A - WOOD IMAGE LIBRARY

TRAINING SET

The photographs presented in this appendix are of the training set of the WOOD library of images that is referred to in this study. For the purposes of display the original images are reduced to one-quarter size by displaying only every second pixel in every second row of each image. In this way four images are displayed on one screen and photographed. Each image is digitised from a section of video tape recorded at the Woods and Forests Department Mount Gambier mill. The size increments in the images are steps of 64 pixels.

Following the photographs is a list of all the features in the training set and the size of the features as determined by a Quality Control Office at the time the images were recorded. The dimensions are in millimetres and the origin is the top left of each image. DX and DY describe the size of a box that encloses the feature.

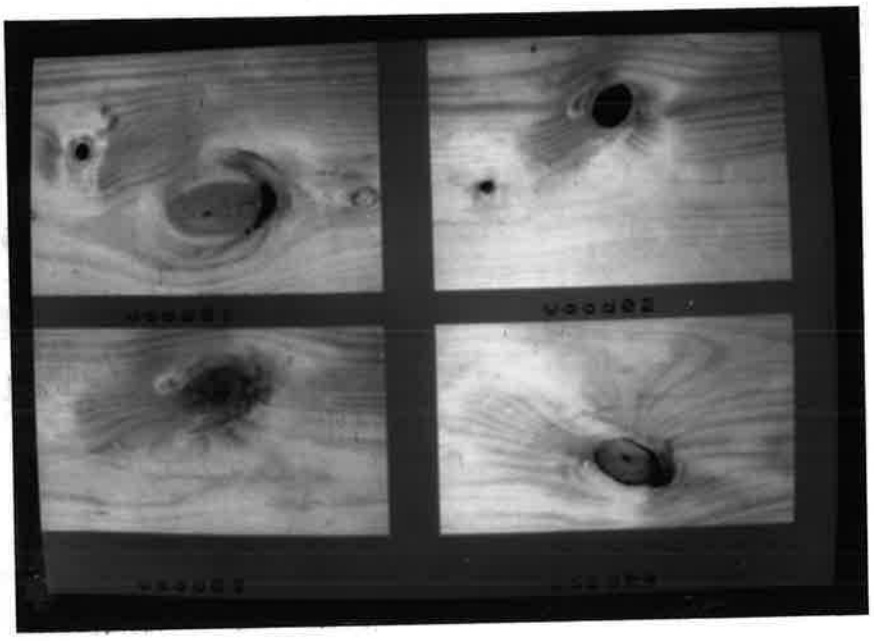

Figure A1A.1 Training set - WOOD01, WOOD02, WOOD03, WOOD04.

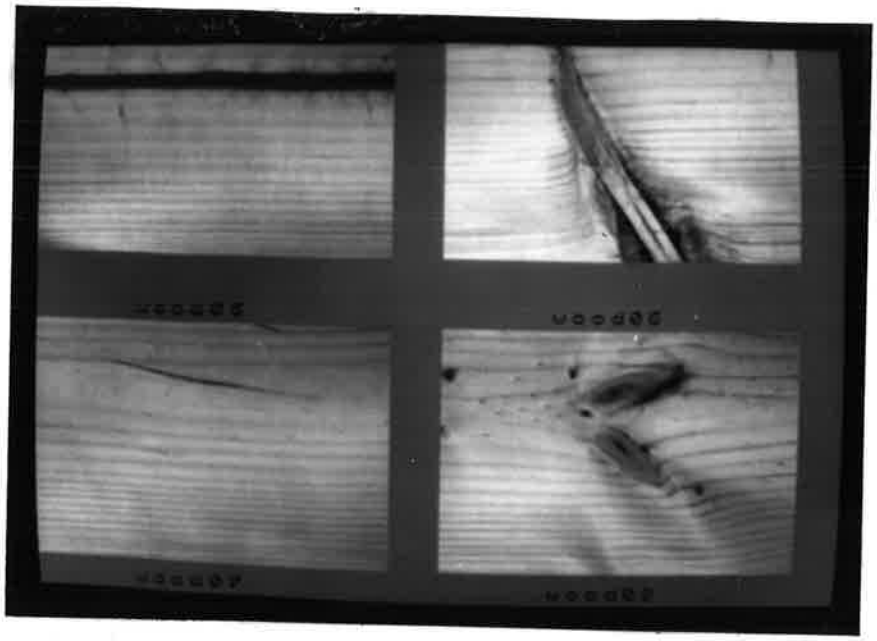

Figure A1A.2 Training set - WOOD05, WOOD06, WOOD07, WOOD08.

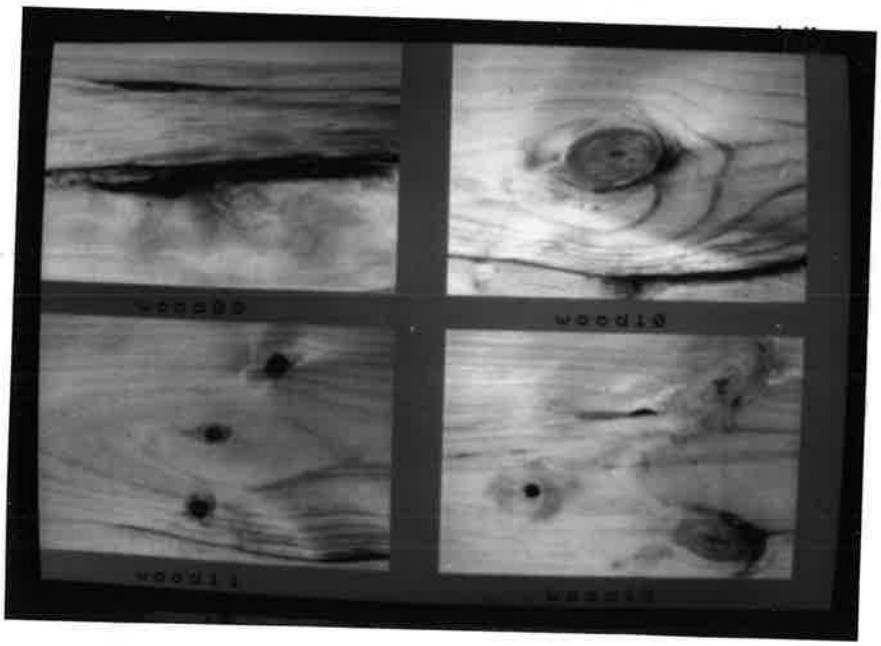

Figure A1A.3 Training set - WOOD09, WOOD10, WOOD11, WOOD12.

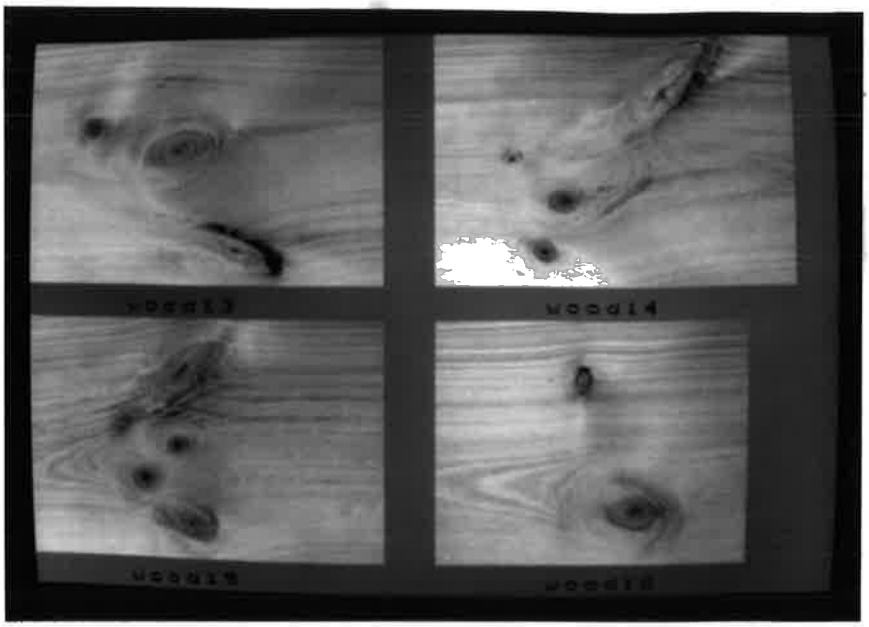

Figure A1A.4 Training set - WOOD13, WOOD14, WOOD15, WOOD16.

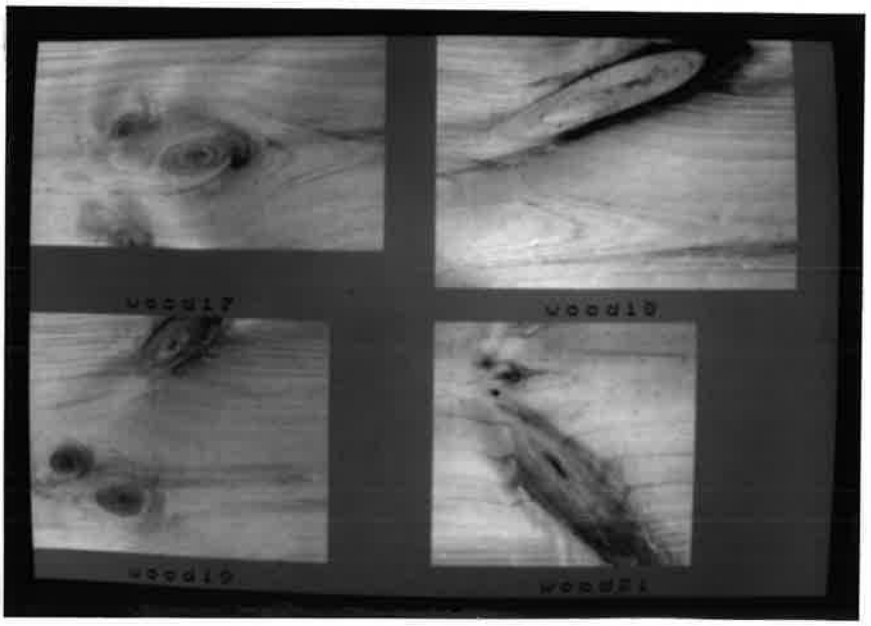

Figure A1A.5 Training set - WOOD17, WOOD18, WOOD19, WOOD21.

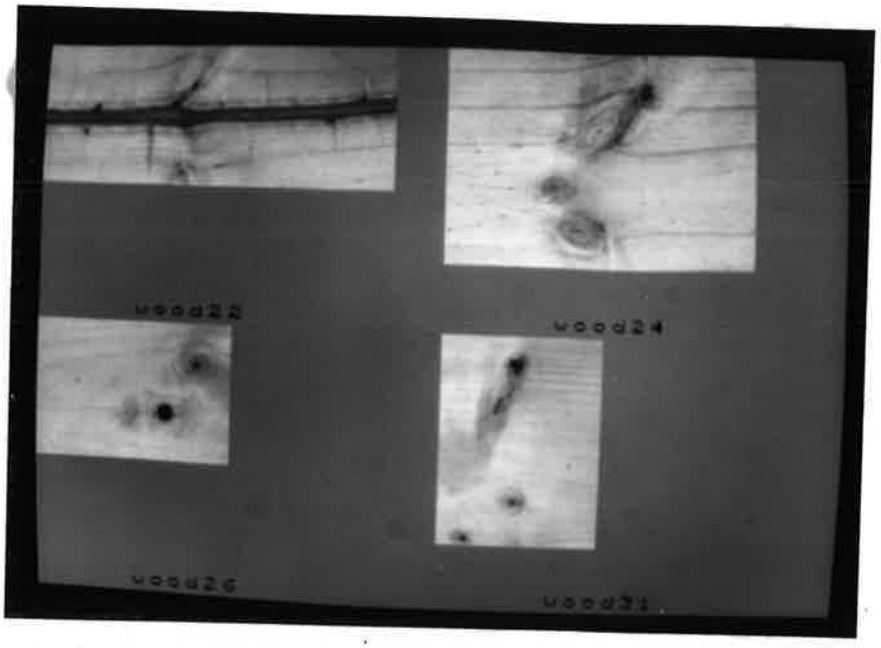

Figure A1A.6 Training set - WOOD22, WOOD24, WOOD26, WOOD31.

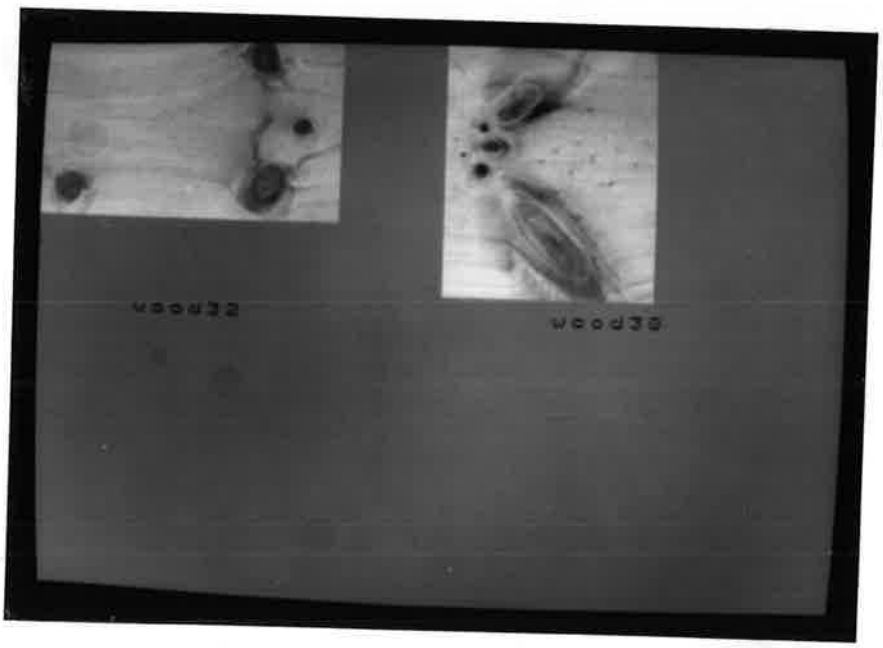

Figure A1A.7 Training set - WOOD32, WOOD38.

TRAINING SET FEATURE LIST

IMAGE	X	Y	DX	DY	TYPE
wood01	34.6	62.0	10.1	11.3	small hole
wood01	100.2	91.0	73.0	37.2	large knot with bark 25%
wood01	156.2	92.5	16.5	31.6	bark 6mm wide
wood02	31.4	96.3	7.5	6.0	pin knot
wood02	106.1	33.1	26.7	28.6	medium hole
wood03	106.6	26.3	42.1	33.5	large knot sound
wood04	101.3	81.2	43.7	33.1	large encased knot
wood05	0.0	18.8	238.8	13.2	pith 10mm
wood06	69.8	0.0	90.1	144.4	large encased spike knot
wood07	55.4	29.0	106.6	18.0	resin
wood07	148.7	0.0	20.8	4.9	resin
wood08	0.5	26.7	7.5	7.9	pin knot
wood08	95.9	21.1	62.9	30.1	large knot with bark 25%
wood08	99.1	64.3	45.8	42.1	largeknot with bark 25%
wood08	81.5	24.8	4.8	7.5	pin knot
wood08	90.6	52.6	4.8	5.6	pin knot
wood08	167.4	105.7	5.9	6.4	pin knot
wood09	29.8	27.4	113.5	7.1	bark pocket
wood09	0.0	75.6	238.8	31.2	bark pocket 20mm
wood10	74.1	53.4	67.7	44.0	large knot with bark 33%
wood10	0.0	138.7	238.8	21.4	bark pocket 4.5mm wide
wood11	157.2	20.7	14.4	16.5	small knot
wood11	119.4	74.1	10.1	9.4	pin knot
wood11	109.8	126.0	11.7	13.5	small knot
wood11	45.8	147.4	192.9	18.0	bark 3mm wide
wood12	53.3	104.9	9.1	8.6	pin knot hole
wood12	118.9	51.1	21.3	5.3	bark
wood12	144.4	1.5	67.2	55.3	large sound knot
wood12	155.1	126.3	60.2	30.8	large knot with bark 33% 5mm
wood13	38.4	53.8	14.4	13.9	small knot
wood13	83.7	62.0	41.6	25.6	medium sound knot
wood13	122.6	126.7	52.8	32.7	large encased knot 50% 8mm bark
wood14	45.8	79.3	6.9	7.1	pin knot
wood14	61.8	138.4	14.4	14.7	small knot
wood14	71.4	103.4	18.7	15.4	small knot
wood14	123.7	11.3	52.8	47.8	large sound knot
wood15	91.7	15.0	33.6	37.2	large sound knot
wood15	97.5	81.6	14.9	12.4	small sound knot
wood15	72.5	105.7	16.0	14.3	small sound knot
wood15	105.5	131.2	24.5	25.2	medium sound knot
wood16	89.0	31.2	9.6	17.3	small encased knot 75%
wood16	109.3	119.9	35.2	24.8	medium sound knot
wood17	56.0	51.5	21.3	18.0	small sound knot
wood17	59.2	131.2	29.8	13.2	small sound knot
wood17	90.6	67.7	61.3	27.1	medium knot with 9mm chip
wood18	25.6	0.0	179.1	78.2	large knot, encased 82% 30mm bark
wood19	15.5	94.4	26.7	19.6	small sound knot
wood19	47.4	120.3	32.5	22.9	medium sound knot

wood19	78.9	0.0	56.5	41.0	large encased knot 25%
wood21	31.4	25.2	4.3	5.3	pin knot
wood21	42.6	28.2	12.3	12.4	small knot with 3.5mm chip
wood21	51.2	61.3	96.5	107.2	large sound knot
wood21	37.8	45.9	3.7	4.1	pin knot
wood22	0.0	33.8	238.8	23.7	pith out of grade > 10mm
wood24	81.0	19.6	51.2	47.4	large bark encased knot 25%
wood24	58.6	83.5	17.6	13.5	small sound knot
wood24	70.9	111.3	31.4	19.9	small sound knot
wood26	87.9	57.2	10.1	11.3	small hole
wood26	106.1	19.2	17.6	16.5	small knot with 3mm chipping
wood31	31.4	15.0	21.3	52.3	large knot with bark 6mm
wood31	14.4	136.5	5.3	5.6	pin knot
wood31	45.3	111.3	4.8	4.9	pin knot
wood32	10.7	91.4	20.3	20.7	medium sound knot
wood32	144.4	0.0	18.1	21.4	medium sound knot
wood32	174.3	51.1	11.7	10.5	small sound knot
wood32	147.6	83.1	22.4	27.1	medium sound knot
wood38	32.5	33.5	41.0	16.5	small knot with bark 25% 6mm
wood38	20.3	51.5	46.4	16.0	pin knot
wood38	24.0	62.8	10.7	8.3	pin knot
wood38	18.7	78.2	8.5	9.4	pin knot
wood38	42.6	90.6	61.3	77.8	large sound knot

APPENDIX 1B - WOOD IMAGE LIBRARY

TEST SET

The photographs presented in this appendix are of the test set of the WOOD library of images that is referred to in this study. For the purposes of display the original images are reduced to one-quarter size by displaying only every second pixel in every second row of each image. In this way four images are displayed on one screen and photographed. Each image is digitised from a section of video tape recorded at the Woods and Forests Department Mount Gambier mill. The size increments in the images are steps of 64 pixels.

Following the photographs is a list of all the features in the test set and the size of the features as determined by a Quality Control Office at the time the images were recorded. The dimensions are in millimetres and the origin is the top left of each image. **DX** and **DY** describe the size of a box that encloses the feature.

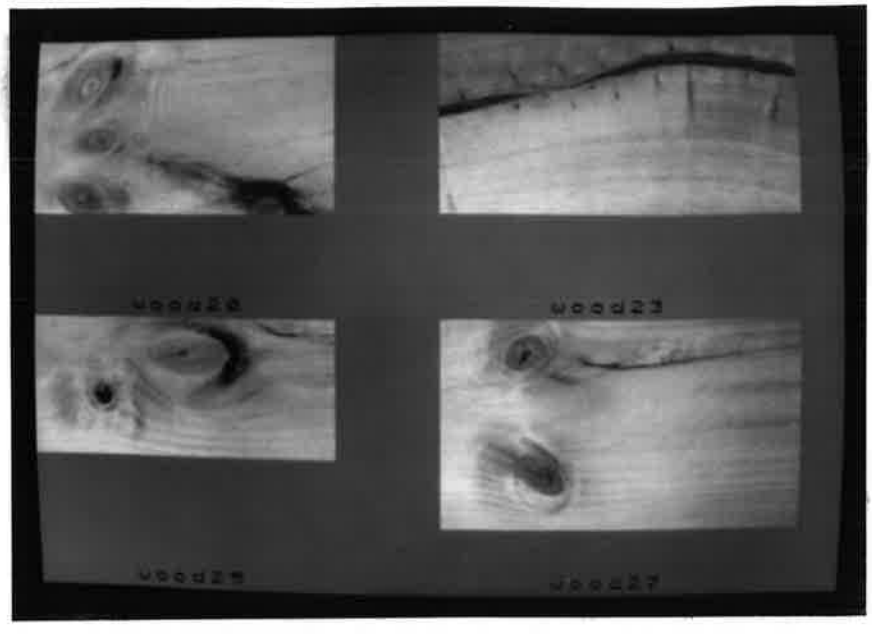

Figure A1B.1 Test set - WOOD20, WOOD23, WOOD25, WOOD27.

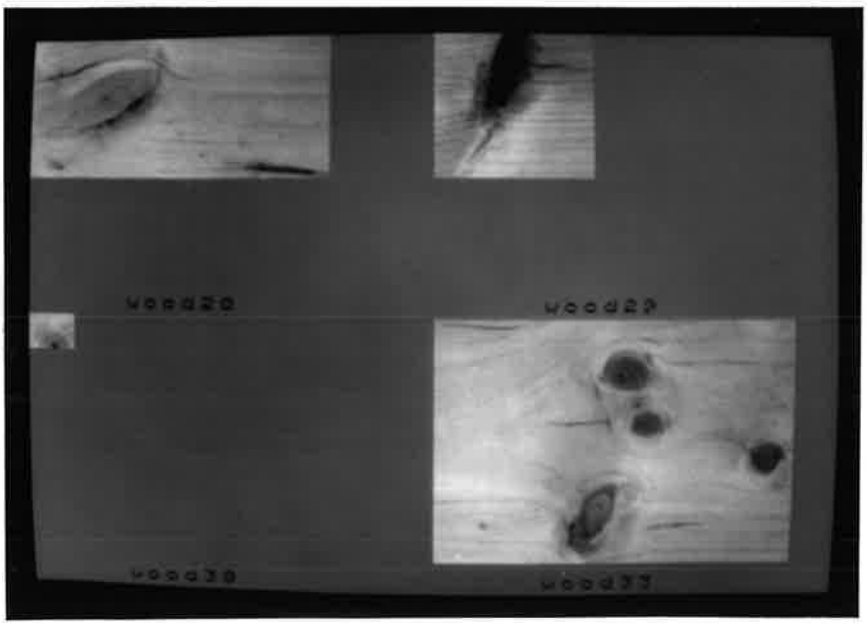

Figure A1B.2 Test set - WOOD28, WOOD29, WOOD30, WOOD33.

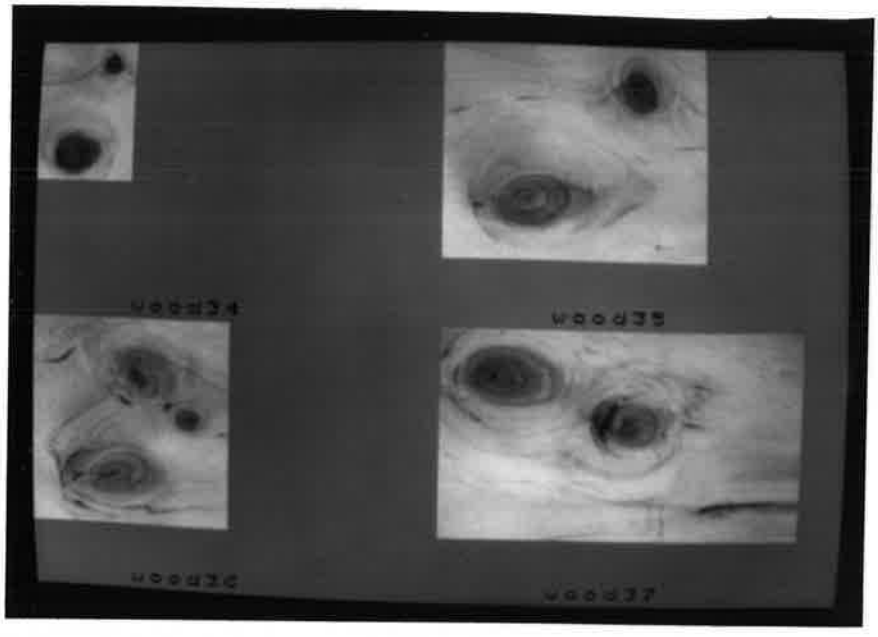

Figure A1B.3 Test set - WOOD34, WOOD35, WOOD36, WOOD37.

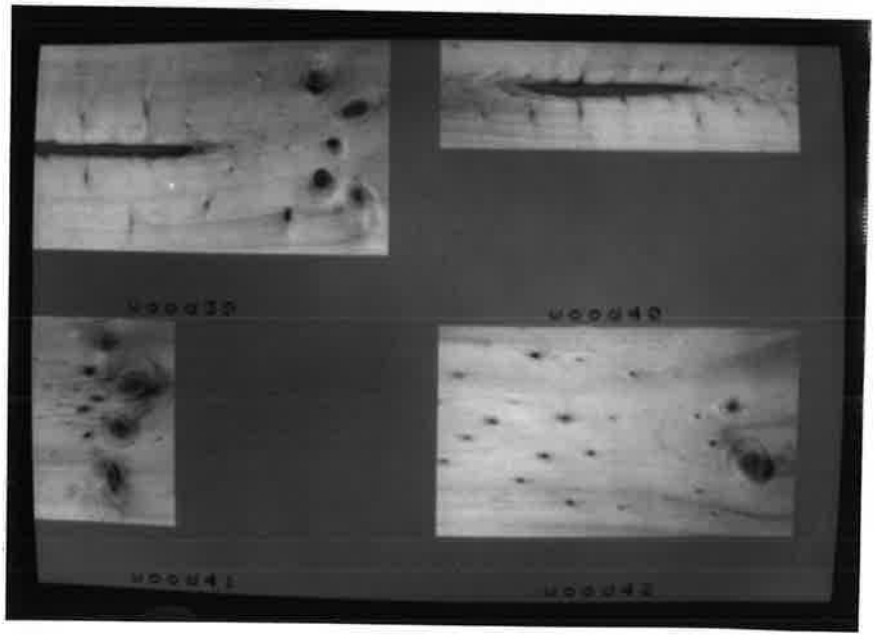

Figure A1B.4 Test set - WOOD39, WOOD40, WOOD41, WOOD42.

Figure A1B.5 Test set - WOOD43, WOOD44, WOOD45, WOOD46.

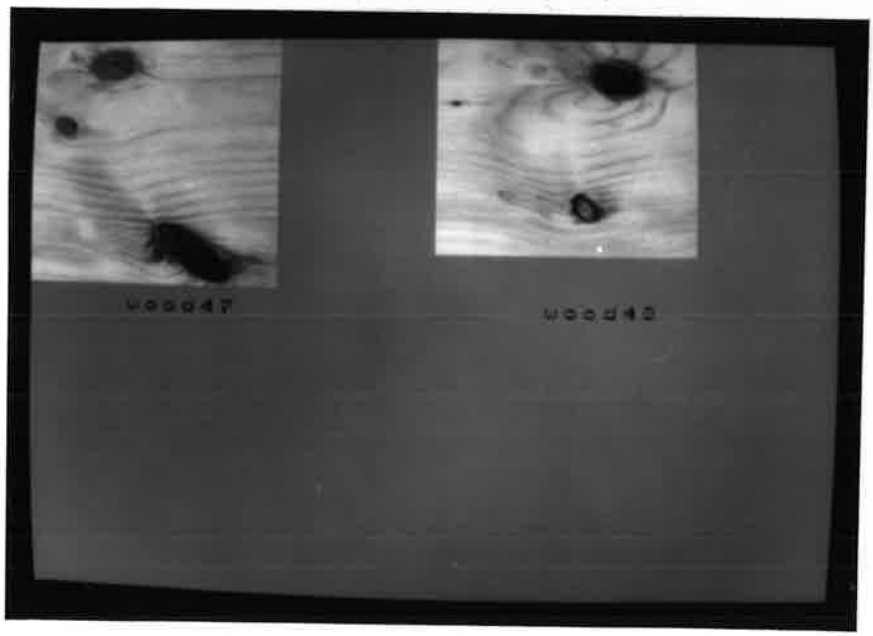

Figure A1B.6 Test set - WOOD47, WOOD48.

TEST SET FEATURE LIST

IMAGE	X	Y	DX	DY	TYPE
wood20	19.7	11.3	41.0	34.2	large knot with chip
wood20	36.8	60.9	20.8	17.3	small sound knot
wood20	18.7	98.9	32.0	21.4	medium sound knot
wood20	138.0	97.4	46.4	22.9	medium knot with bark 75% 8mm
wood23	0.0	39.9	55.4	16.5	pith
wood23	95.9	11.3	142.8	20.7	pith
wood25	45.3	47.8	10.1	13.2	small hole
wood25	87.4	12.4	45.8	26.3	medium sound knot
wood25	115.7	6.4	26.7	38.7	bark for above knot 7mm wide
wood27	41.0	10.2	24.5	24.1	medium knot with chipping 3mm
wood27	46.9	91.4	32.0	29.0	medium sound knot
wood27	89.5	23.7	100.2	9.8	resin pocket
wood28	34.1	21.1	47.4	40.2	large encased knot 25% 4mm
wood28	9.6	12.8	98.6	15.4	resin
wood28	141.2	87.2	54.4	7.1	resin and bark
wood29	31.4	0.0	26.7	51.1	large hole
wood30	8.0	12.0	22.9	12.0	small sound knot
wood33	109.3	24.1	28.2	24.1	medium sound knot
wood33	127.4	62.8	19.7	18.0	small sound knot
wood33	86.9	115.1	30.4	42.1	large encased knot 25% 8mm
wood33	211.1	85.4	21.3	19.2	small encased knot 25%
wood33	4.3	2.3	50.6	4.5	resin
wood33	68.2	68.4	43.7	3.4	resin
wood33	136.4	140.6	40.5	4.5	resin
wood34	45.8	7.5	10.7	14.3	small encased knot 25%
wood34	13.3	63.5	33.0	27.8	medium sound knot
wood35	35.2	87.2	45.3	35.0	large sound knot
wood35	114.1	15.8	22.4	28.6	medium sound knot
wood36	68.2	28.6	21.3	26.7	medium sound knot
wood36	101.3	62.0	16.5	14.7	small sound knot
wood36	42.1	93.6	38.9	29.7	medium sound knot
wood36	19.2	94.8	14.9	27.8	bark encasement
wood37	15.5	11.7	62.4	40.2	large sound knot
wood37	98.1	50.8	40.5	26.3	medium encased knot 25% 12mm
wood37	166.3	116.9	61.8	7.1	resin
wood39	0.0	67.7	120.5	13.2	pith 9mm
wood39	186.0	20.3	11.2	15.0	small encased knot 25%
wood39	208.4	41.7	17.6	10.2	small sound knot
wood39	200.4	66.2	7.5	7.5	pin knot
wood39	191.9	87.2	11.2	12.8	small knot
wood39	174.3	114.3	2.7	8.3	pin knot
wood39	216.4	99.3	6.9	11.3	small sound knot
wood40	49.0	23.7	127.4	14.3	pith 9mm
wood41	51.2	9.0	8.5	10.2	small knot
wood41	40.5	34.2	5.9	6.8	pin knot
wood41	69.3	42.1	22.9	12.8	small knot with bark 5mm
wood41	43.7	54.1	9.1	4.5	pin knot
wood41	37.3	62.4	5.3	3.8	pin knot

wood41	56.5	70.3	14.9	13.2	small sound knot
wood41	40.0	80.8	3.2	4.1	pin knot
wood41	54.4	103.8	10.1	21.1	medium encased knot 25% 1.5mm
wood42	189.7	51.5	5.3	5.6	pin knot
wood42	2.1	90.6	3.7	2.6	needle trace
wood42	17.6	74.4	3.2	2.3	needle trace
wood42	33.0	62.8	3.2	2.3	needle trace
wood42	51.7	103.4	3.2	3.0	needle trace
wood42	59.7	16.9	4.3	4.1	needle trace
wood42	67.2	86.1	2.7	3.4	needle trace
wood42	80.0	59.8	2.1	3.0	needle trace
wood42	88.5	20.3	2.7	2.3	needle trace
wood42	84.2	118.4	2.7	2.6	needle trace
wood42	97.0	83.8	3.7	3.8	needle trace
wood42	123.1	30.1	2.7	2.6	needle trace
wood42	117.8	122.6	2.1	2.3	needle trace
wood42	167.9	59.4	2.1	2.3	needle trace
wood42	179.1	77.1	3.2	3.8	needle trace
wood42	170.0	111.3	2.7	3.0	needle trace
wood42	200.4	86.1	22.9	17.3	small bark encased knot 25% 3mm
wood43	32.5	0.0	72.0	55.3	large sound knot
wood43	14.9	72.6	22.4	17.3	small sound knot
wood43	32.0	104.9	29.3	28.2	medium sound knot
wood44	20.8	34.6	8.0	8.3	pin knot
wood44	50.1	6.4	23.5	16.2	small sound knot
wood44	76.2	132.7	26.7	35.7	large encased knot 100%
wood45	17.1	53.8	17.6	14.3	small sound knot
wood45	102.3	3.8	75.7	32.7	large encased knot 25% 2.5mm
wood46	34.6	21.1	72.0	51.1	large encased knot 100% 7mm
wood47	41.6	6.0	28.2	20.7	medium sound knot
wood47	17.1	53.4	16.5	12.8	small sound knot
wood47	93.3	126.0	45.3	39.9	large hole
wood48	9.1	41.7	4.3	3.0	pin knot
wood48	102.9	15.0	26.1	20.3	medium hole
wood48	88.5	104.9	15.5	15.8	small encased knot 100% 2mm

APPENDIX 2

VISUALISATION OF 4-D TONAL PATTERN SPACE

The photographs in the following pages present the 4-dimensional tonal pattern space in each of the six 2-dimensional viewing planes. Both the training set and test set are presented on the same page for comparison. The samples are coloured according to the distance of the sample from the hyperplane of equation 9.1a:

$$w = (127.97, -140.58, 101.23, -139.05, 5.54)$$

where each sample is of the form:

$$x = (\text{mean}, \text{std. dev.}, \text{skewness}, \text{kurtosis}, 1)$$

The colour scale indicates the relative distance of each sample from the hyperplane that separates the feature measure space into the two sections of the dichotomy. The two classes are represented by circles for clear local areas and crosses for feature local areas. Samples that are correctly classified are blue turning to green as they near the decision boundary. Samples that are misclassified by the decision boundary are coloured red and turn yellow to white as they move further from the boundary. Red circles are clear areas that are wrongly labelled as feature areas (false positives) and red crosses are feature areas that are wrongly labelled as clear areas (false negatives).

The display is scaled in each case so that all the samples in the set fill the screen. This gives a slight linear distortion when comparing the training and the test set in the same space but the colour scale is an invariant characteristic in each photograph.

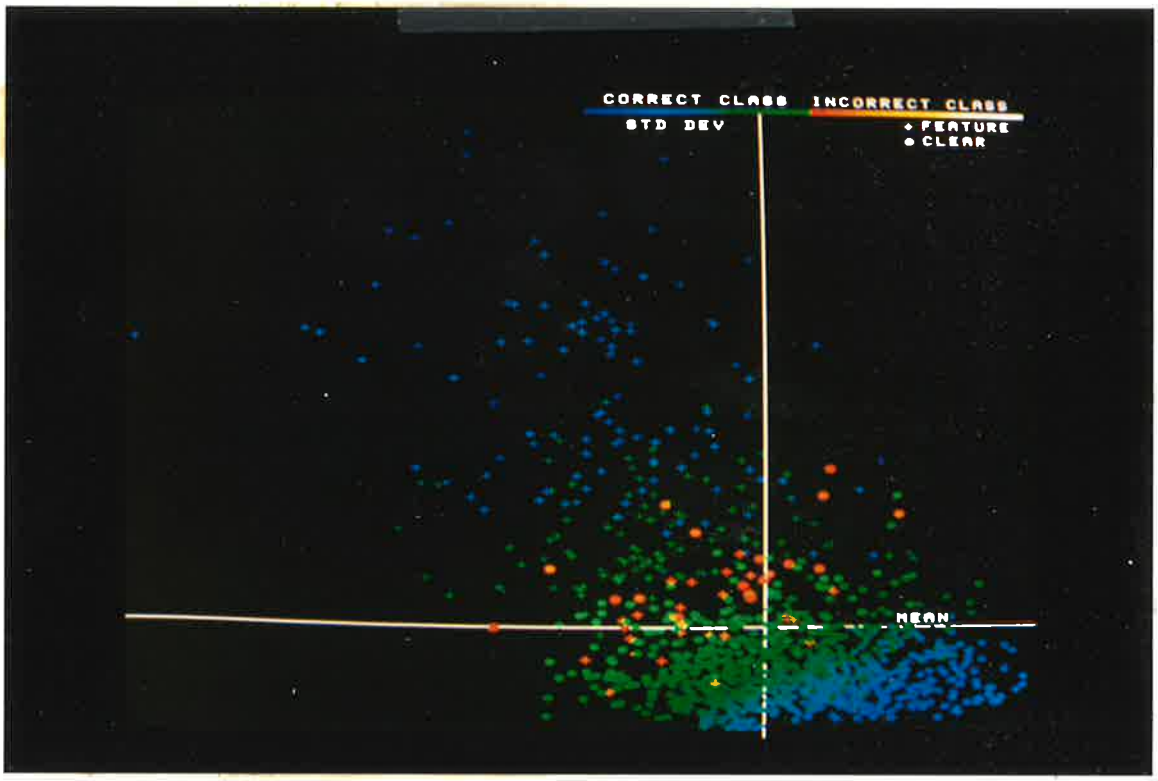

Figure A2.1a 4-d pattern space for the WOOD training set. Mean vs. Standard Deviation.

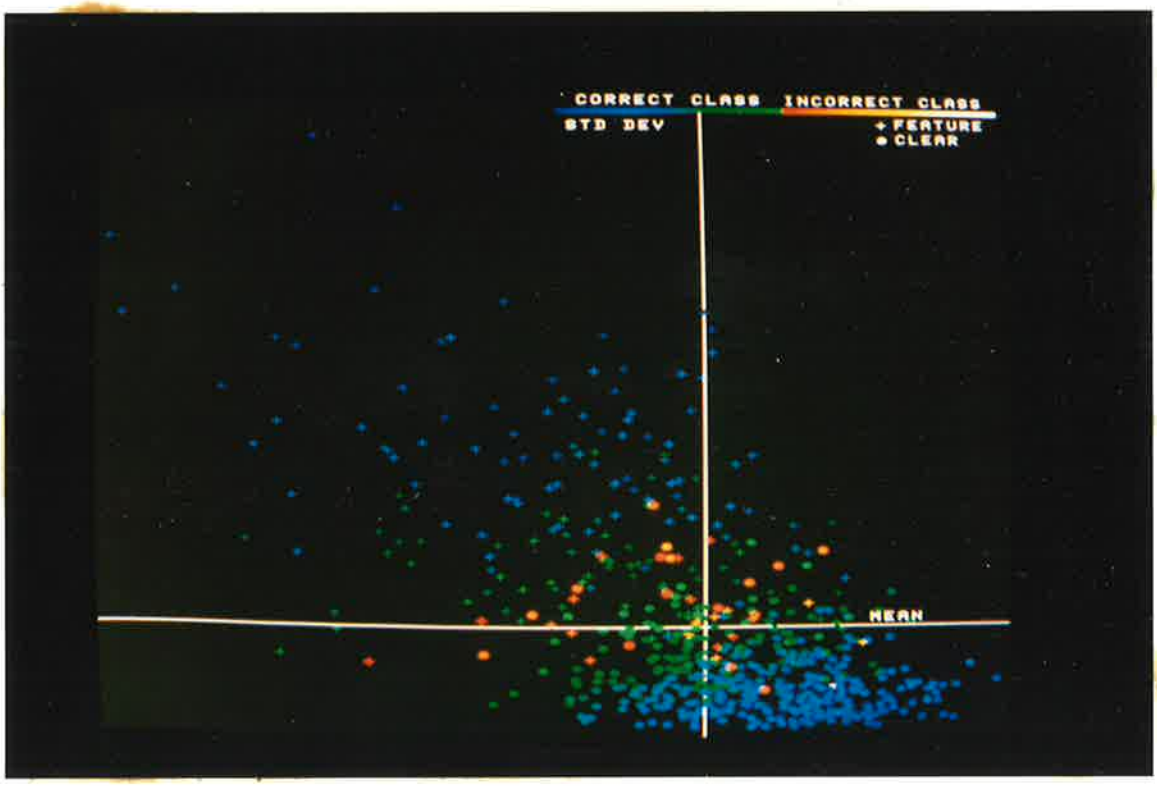

Figure A2.1b 4-d Pattern space for the WOOD test set. Mean vs. Standard Deviation.

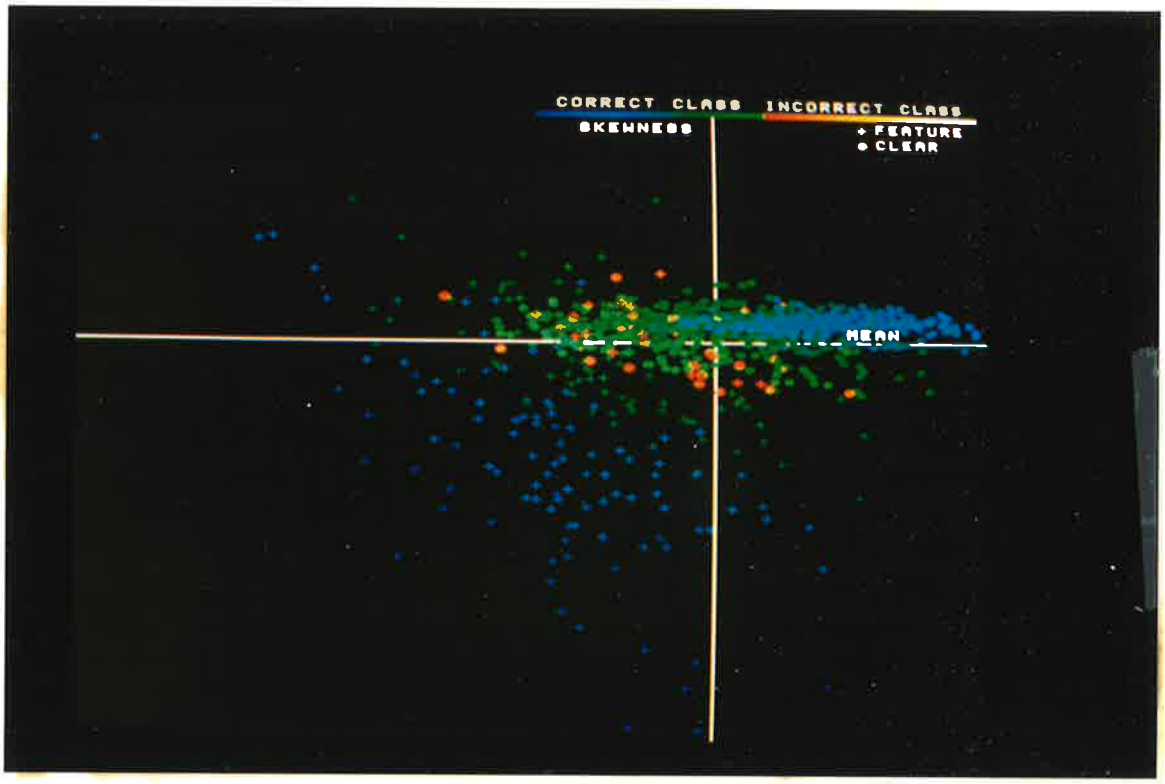

Figure A2.2a 4-d pattern space for the WOOD training set. Mean vs. Skewness.

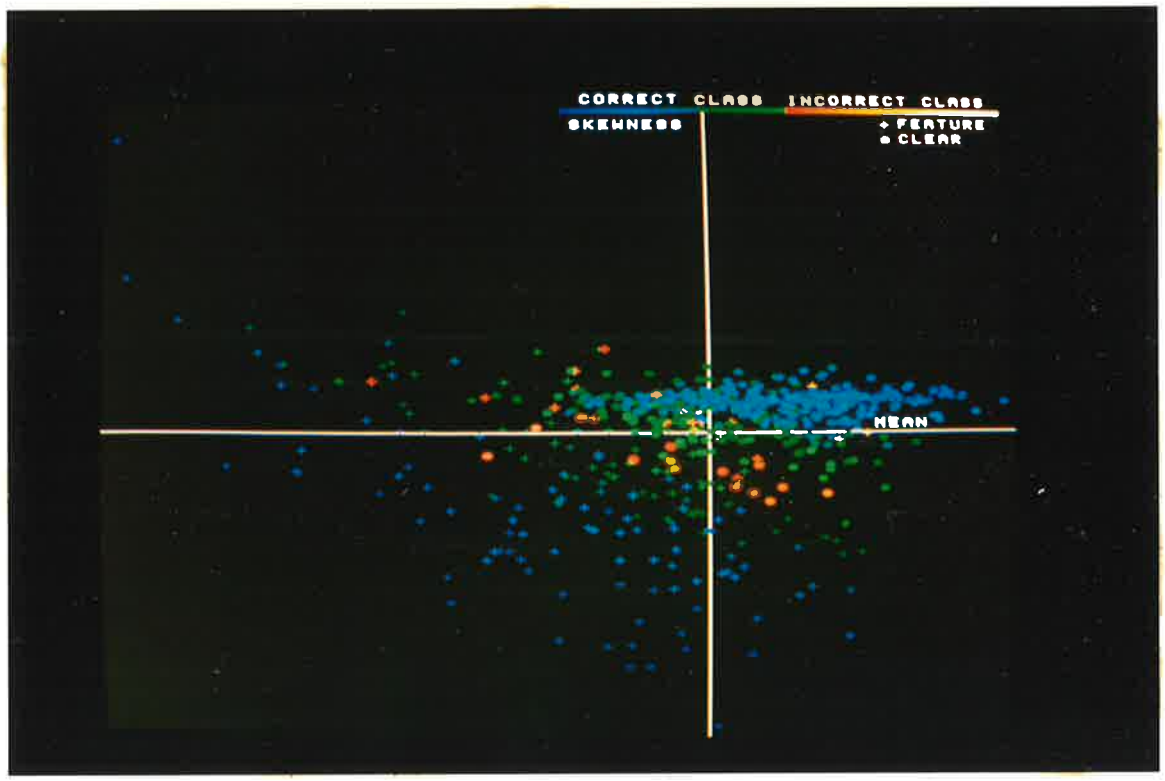

Figure A2.2b 4-d Pattern space for the WOOD test set. Mean vs. Skewness.

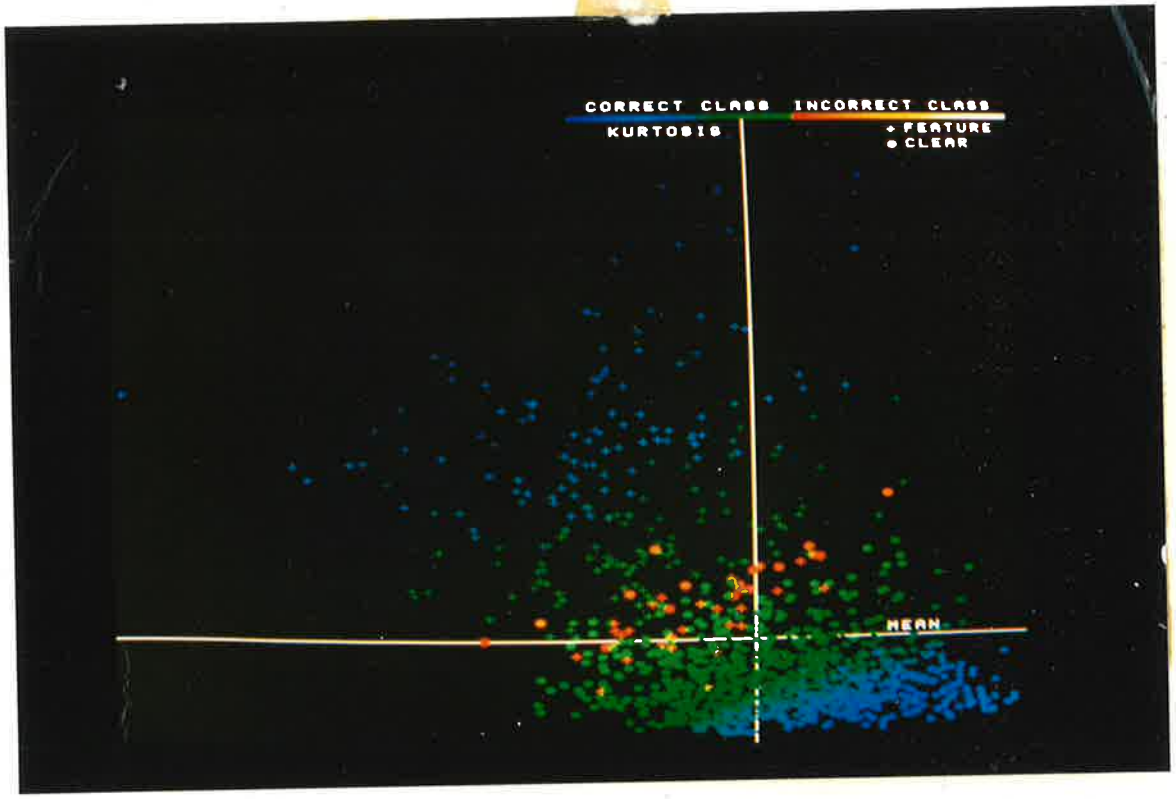

Figure A2.3a 4-d pattern space for the WOOD training set. Mean vs. Kurtosis.

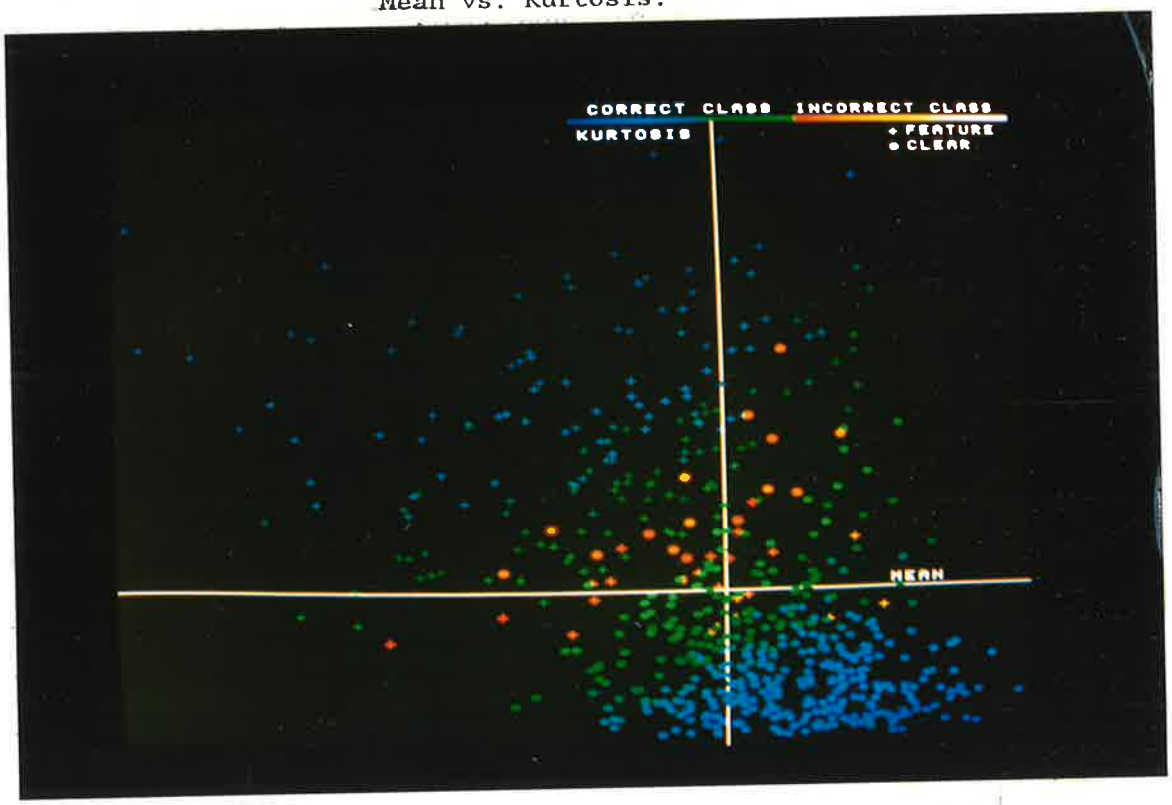

Figure A2.3b 4-d Pattern space for the WOOD test set. Mean vs. Kurtosis.

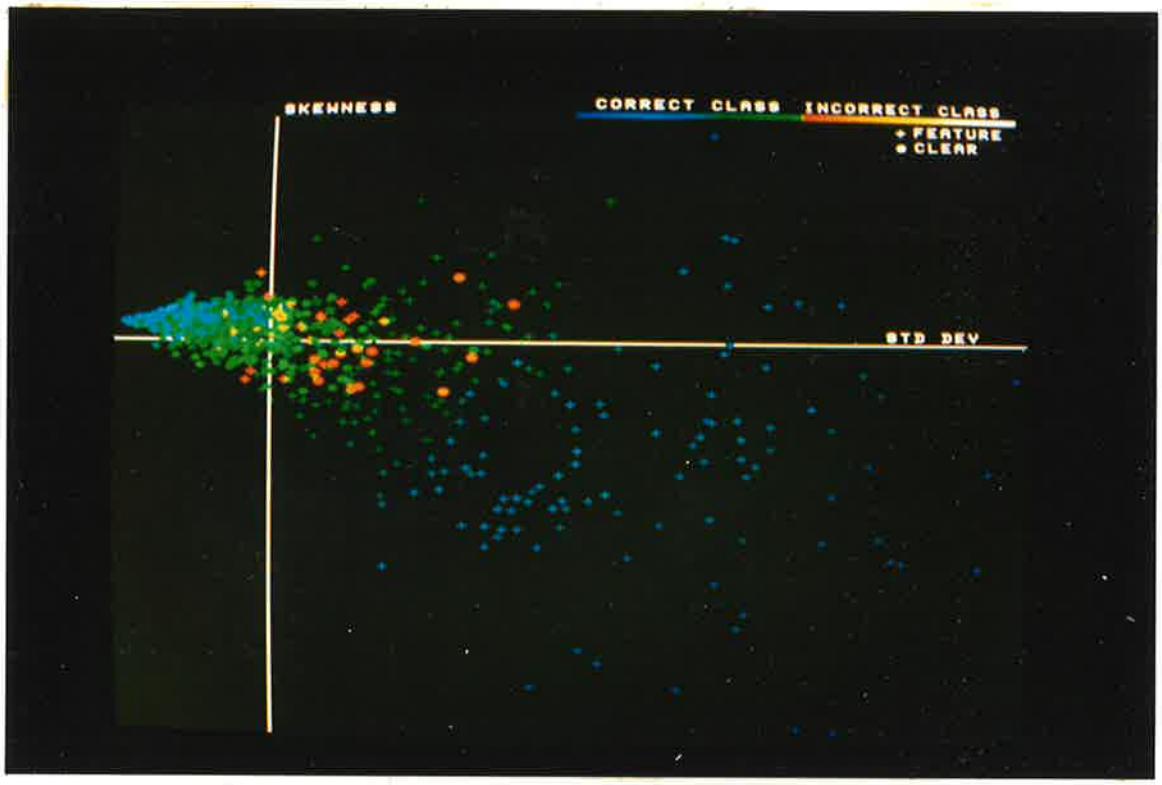

Figure A2.4a 4-d pattern space for the WOOD training set. Standard Deviation vs. Skewness.

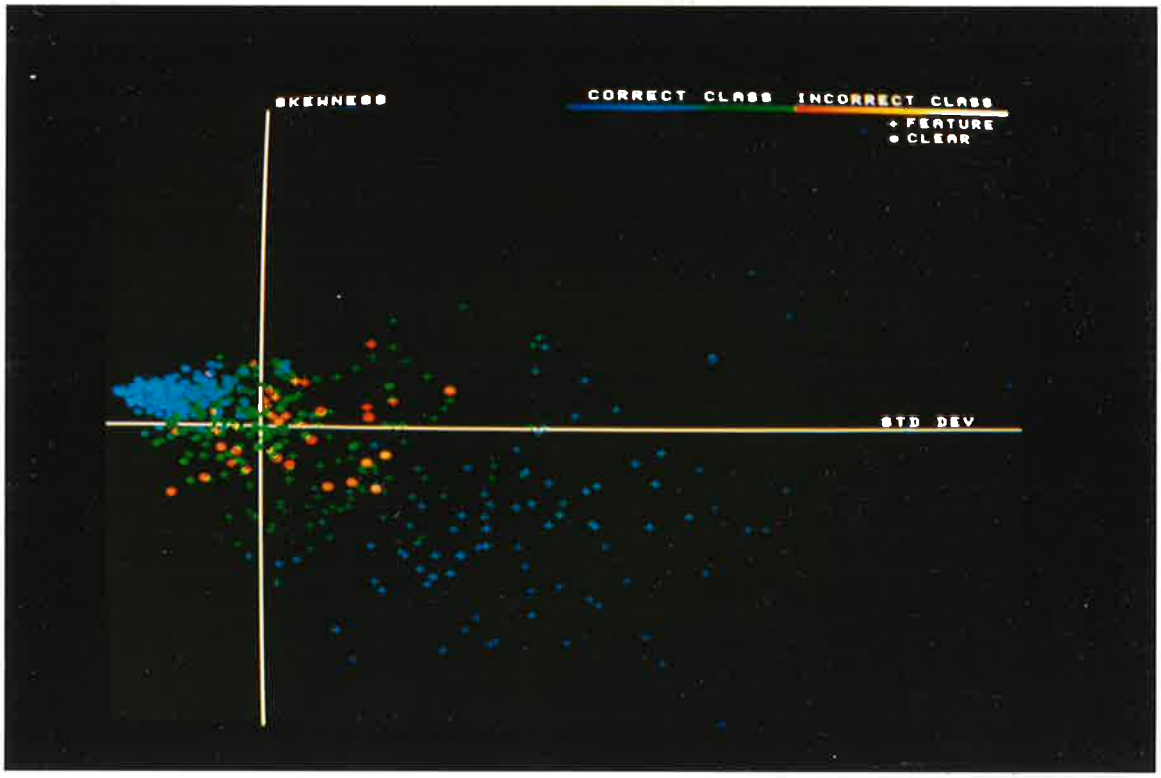

Figure A2.4b 4-d Pattern space for the WOOD test set. Standard Deviation vs. Skewness.

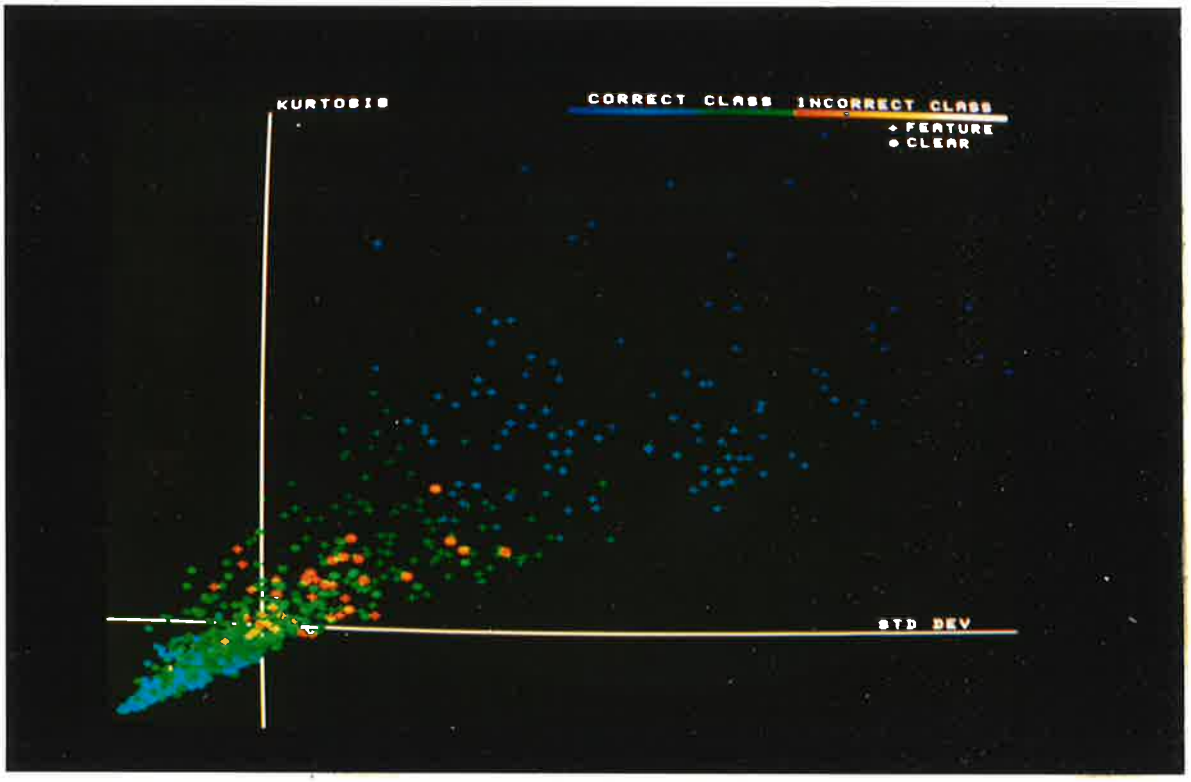

Figure A2.5a 4-d pattern space for the WOOD training set. Standard Deviation vs. Kurtosis.

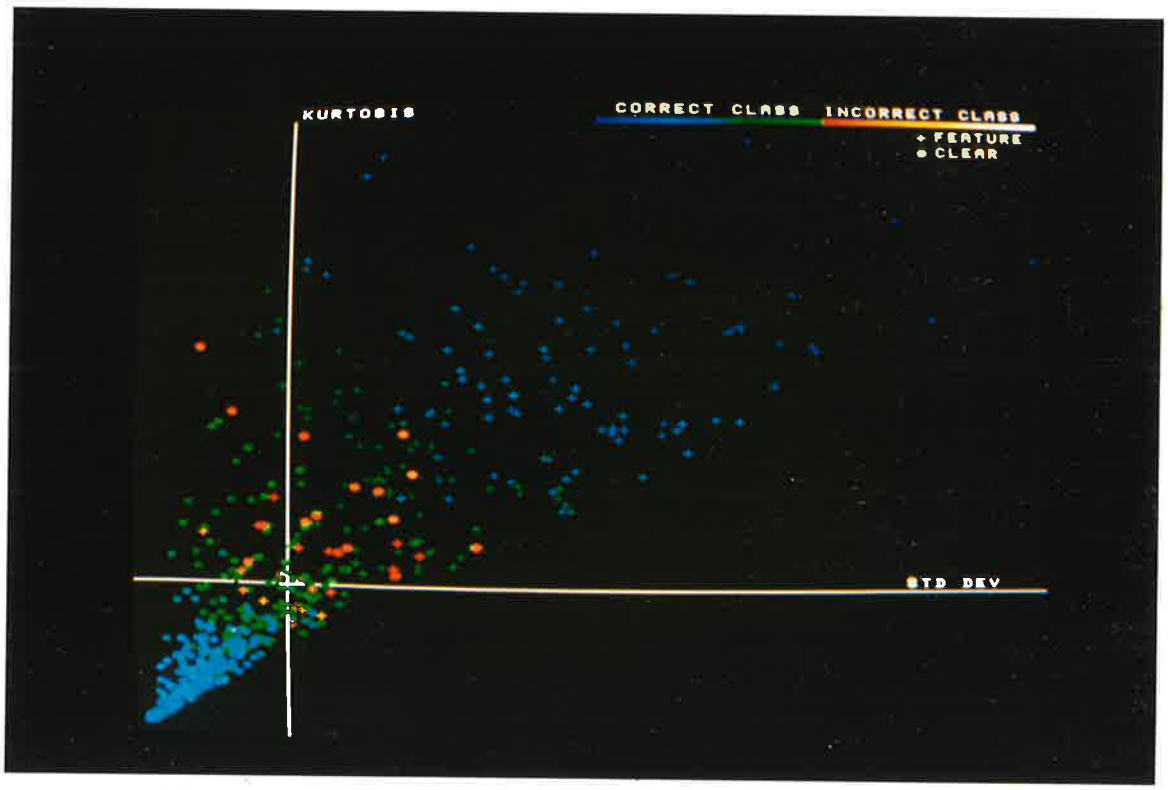

Figure A2.5b 4-d Pattern space for the WOOD test set. Standard Deviation vs. Kurtosis.

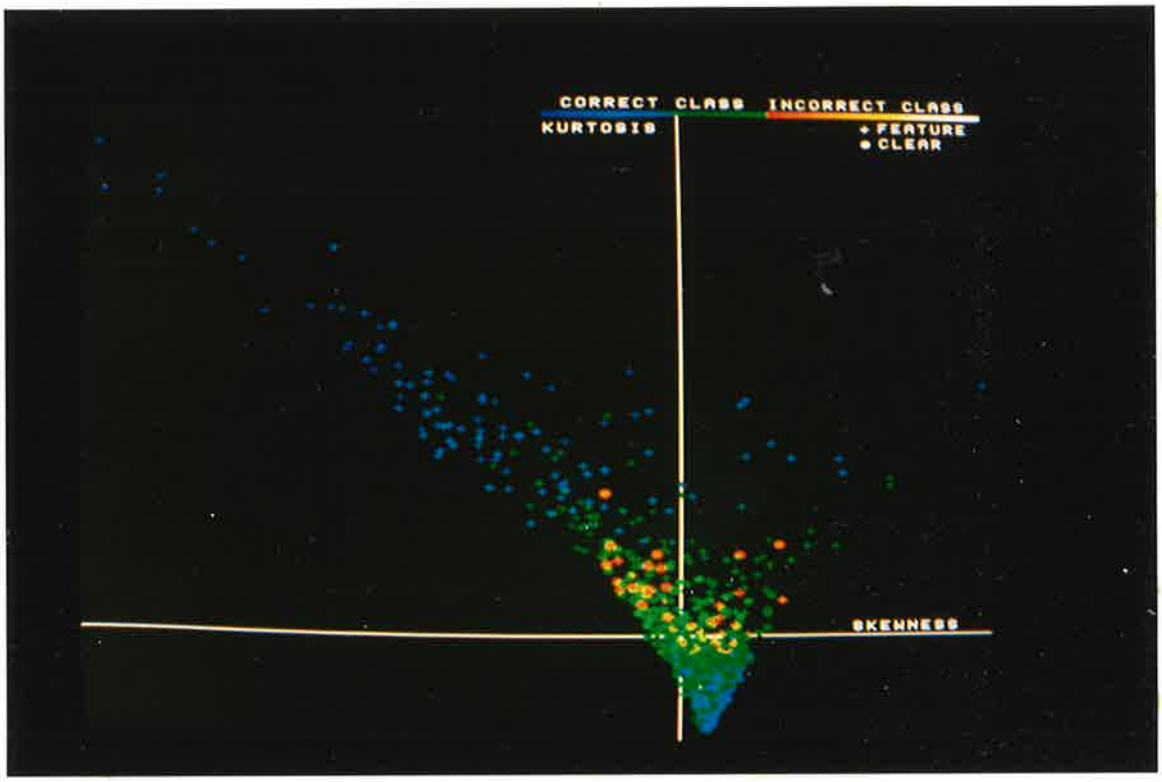

Figure A2.6a 4-d pattern space for the WOOD training set. Skewness vs. Kurtosis.

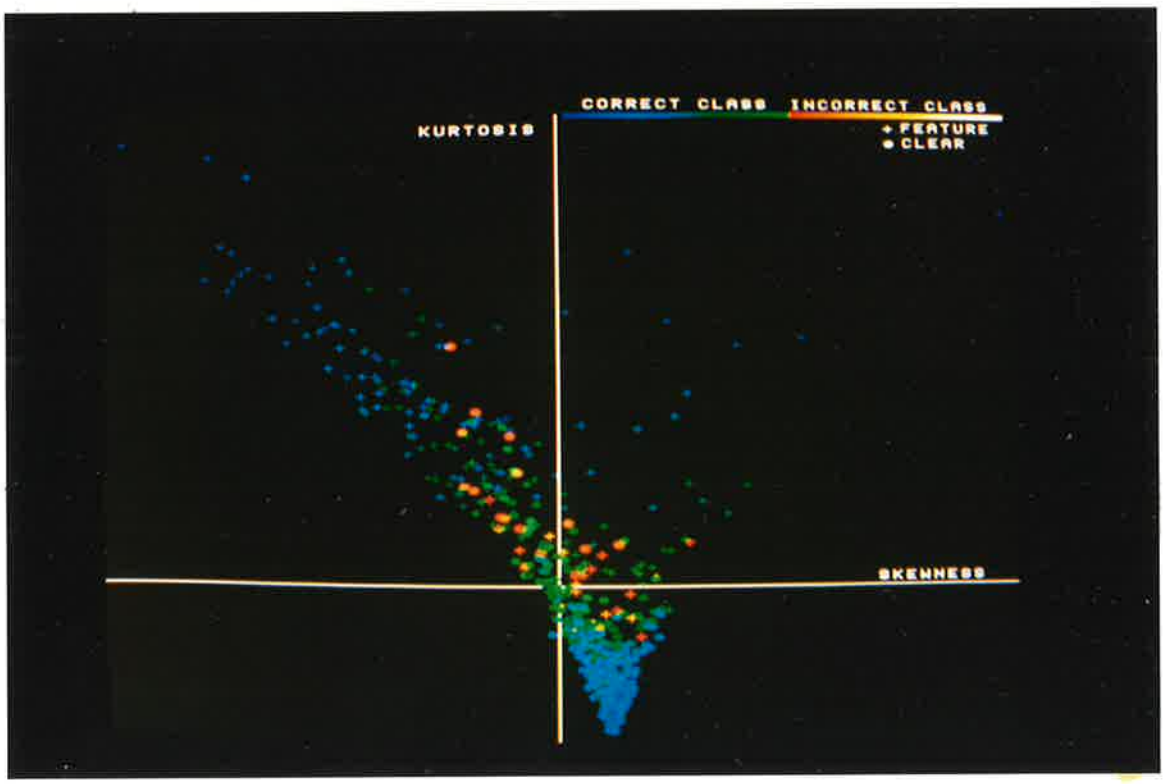

Figure A2.6b 4-d Pattern space for the WOOD test set. Skewness vs. Kurtosis.

APPENDIX 3A - FEATURE DISCRIMINATION

TRAINING SET - LINE STATISTICS METHOD

The photographs presented in this appendix are the processed results of the training set of the WOOD library of images. For the purposes of display the original images are reduced to one-quarter size by displaying only every second pixel in every second row of each image. In this way four images are displayed on one screen and photographed.

Each figure is the result of processing using the line statistics method to determine the feature area extent. Each feature area is processed using the modification of the adaptive threshold method for feature areas. The grid lines are the boundaries of the local areas of size 64 * 64 pixels. The white rectangles are the feature areas and each coloured blob is separate from every other. Vacant areas that can be inferred by the absence of grid lines are feature areas that have been detected and classed as clear.

Following the figures is a list of the location and size of every blob with various descriptive shape measures. This is followed by a description of how well each feature is described.

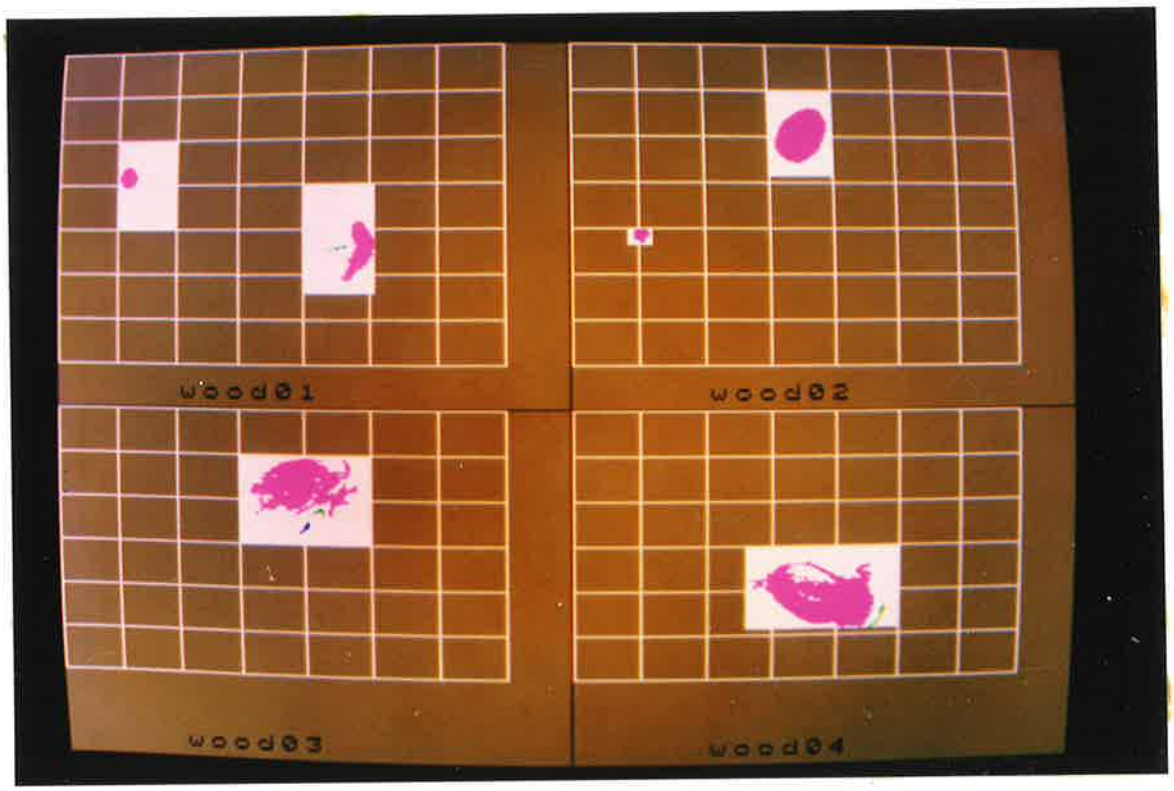

Figure A3A.1 Training set - WOOD01, WOOD02, WOOD03, WOOD04.

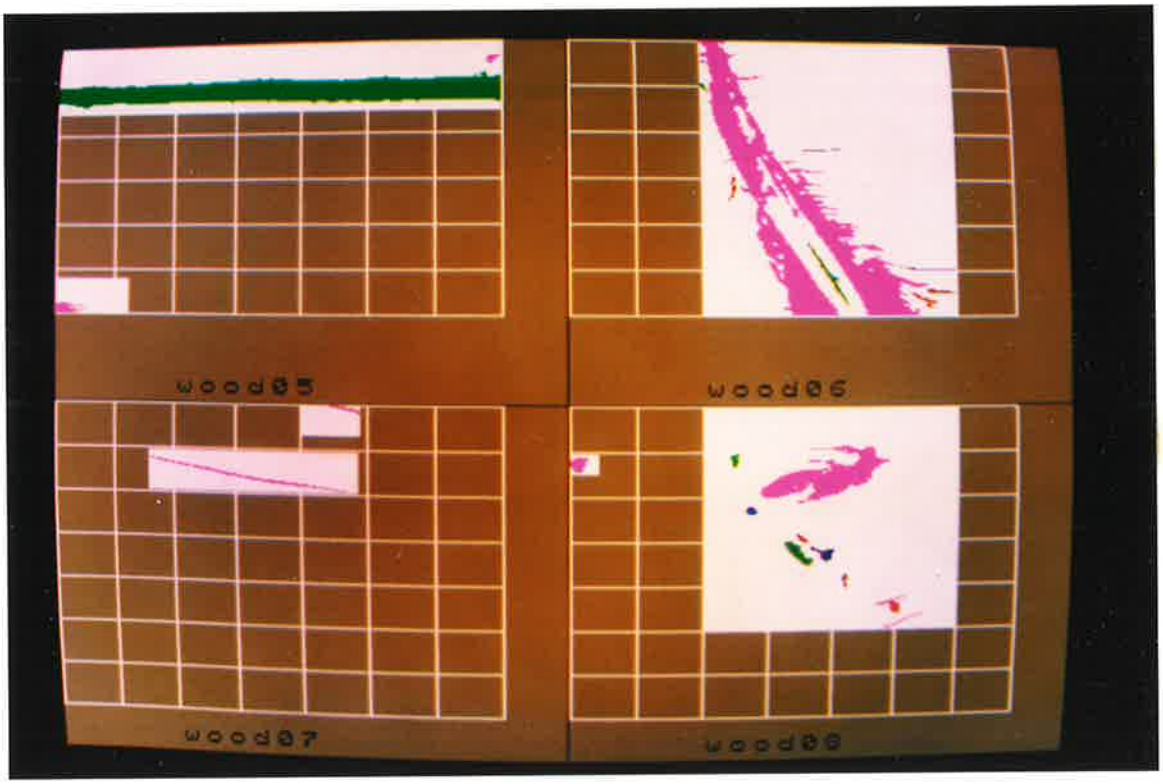

Figure A3A.2 Training set - WOOD05, WOOD06, WOOD07, WOOD08.

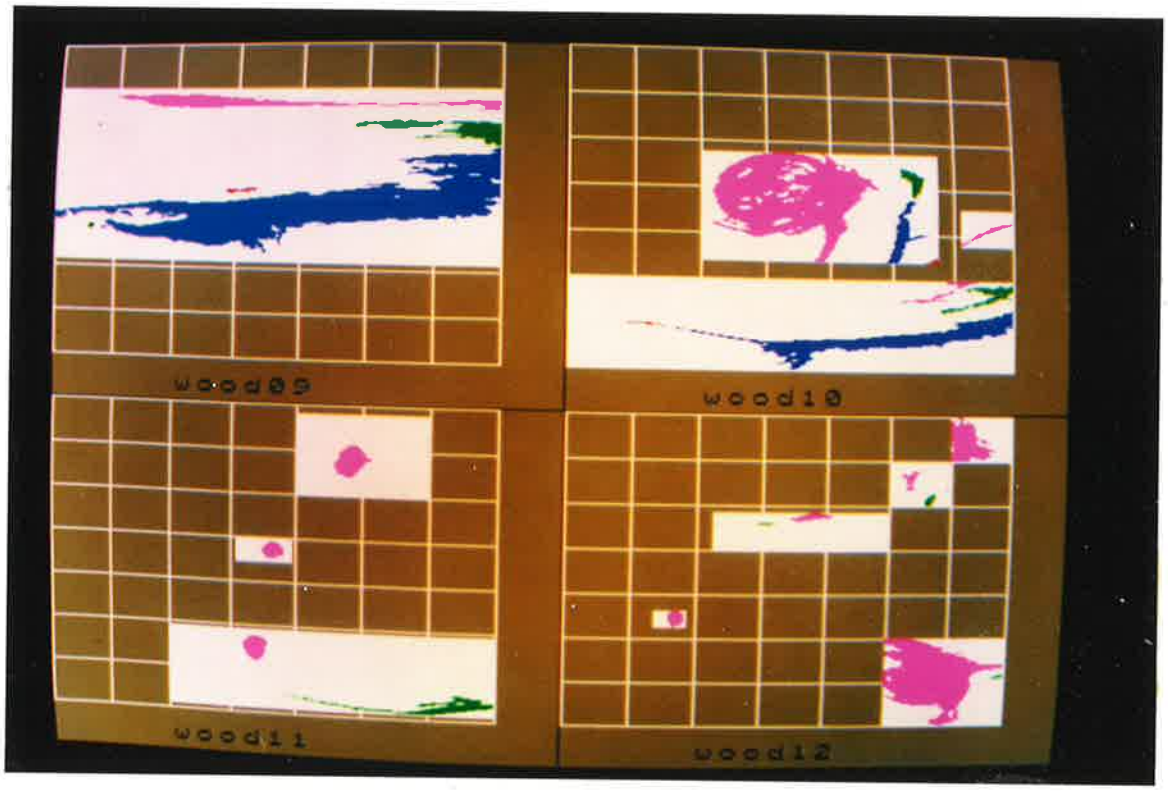

Figure A3A.3 Training set - WOOD09, WOOD10, WOOD11, WOOD12.

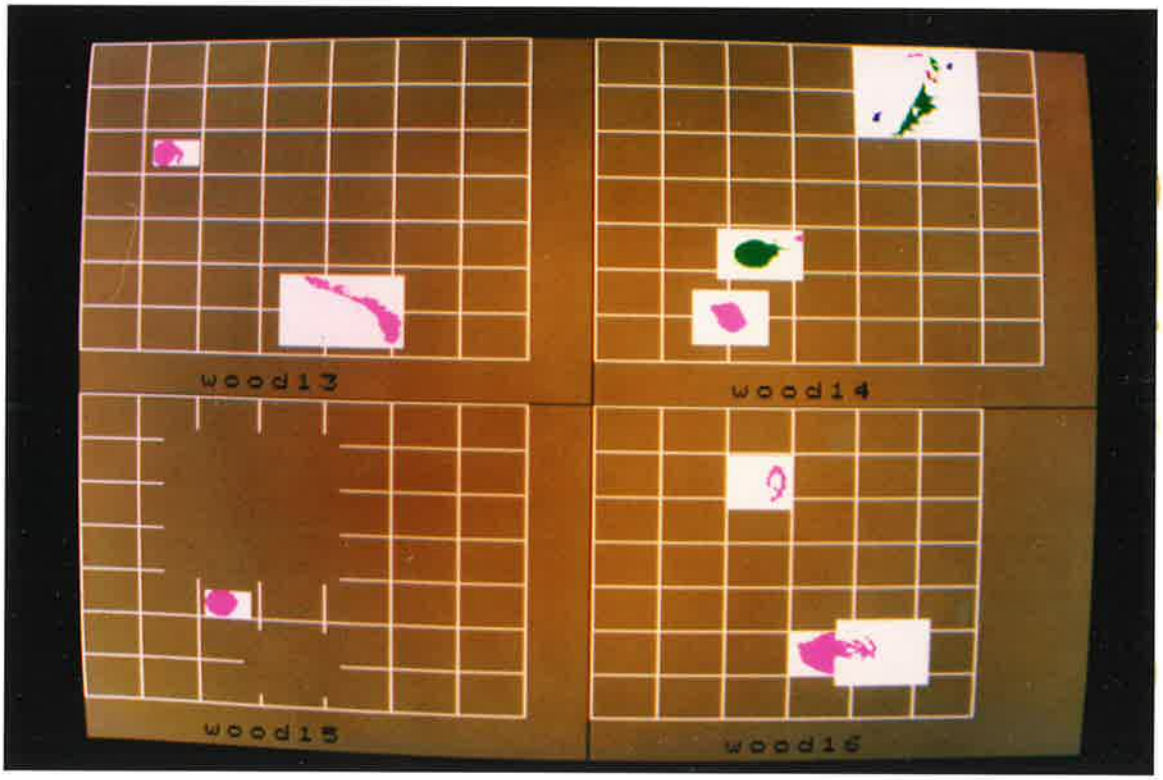

Figure A3A.4 Training set - WOOD13, WOOD14, WOOD15, WOOD16.

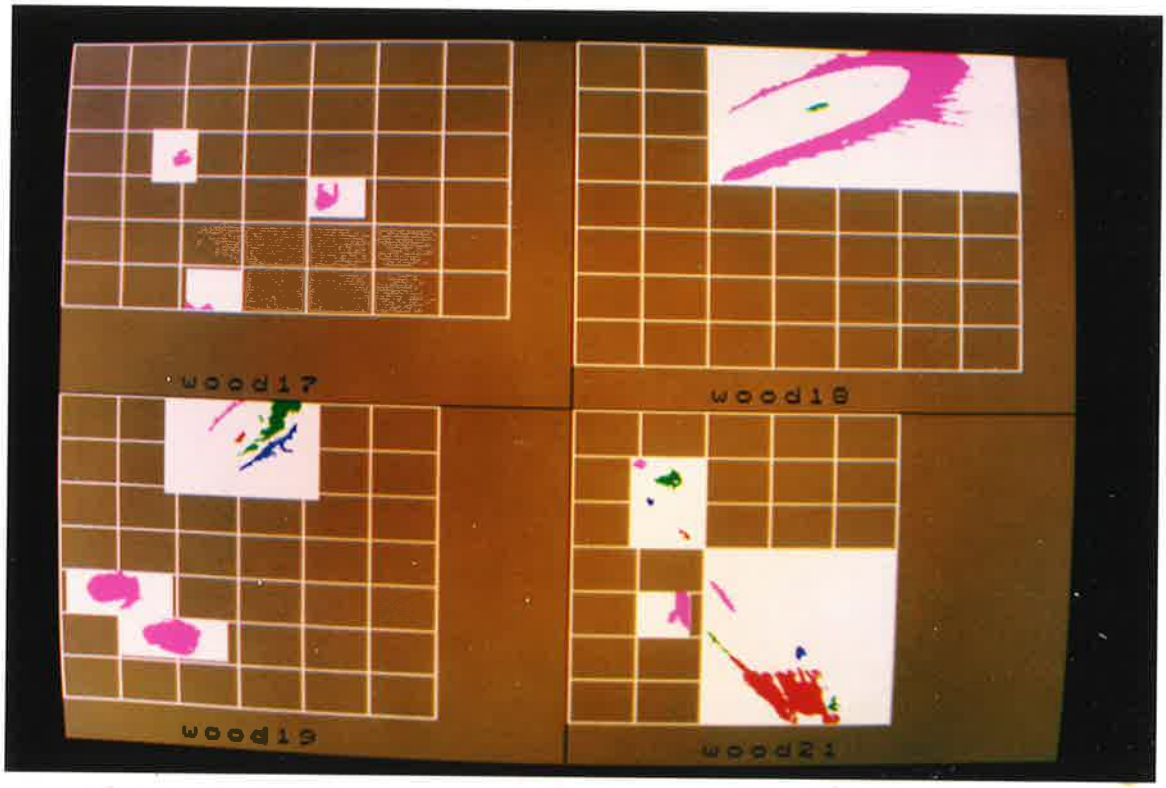

Figure A3A.5 Training set - WOOD17, WOOD18, WOOD19, WOOD21.

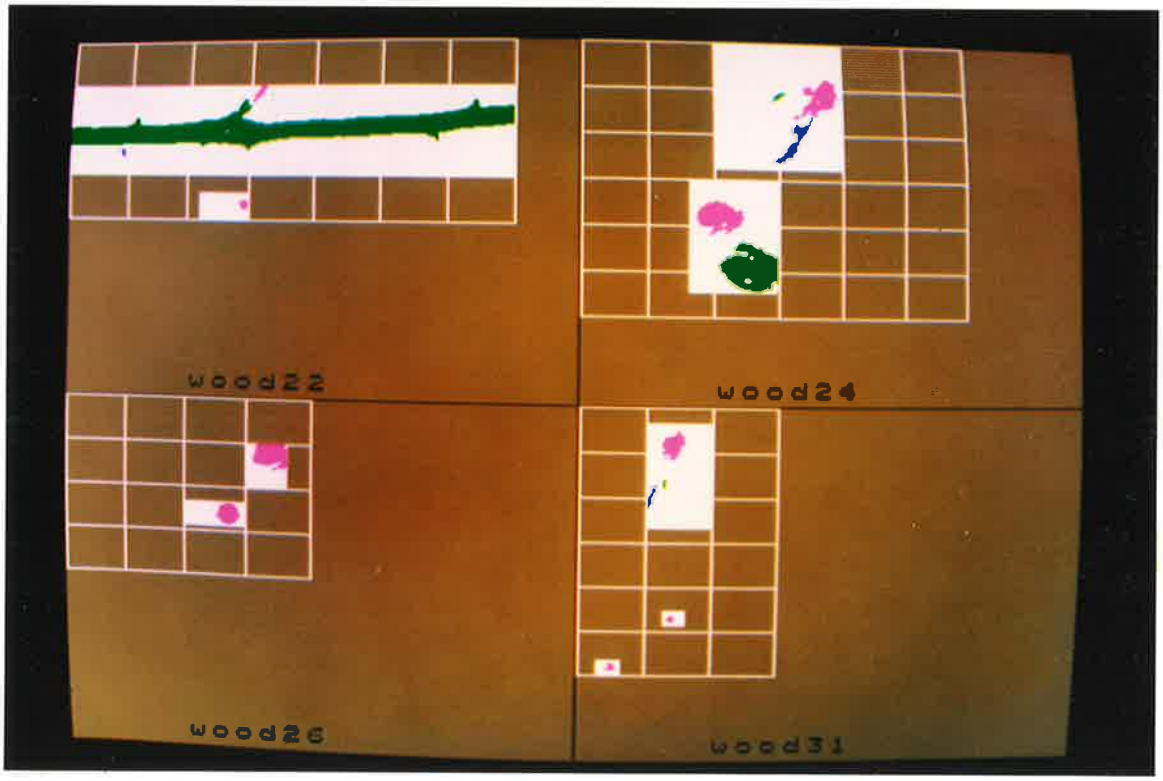

Figure A3A.6 Training set - WOOD22, WOOD24, WOOD26, WOOD31.

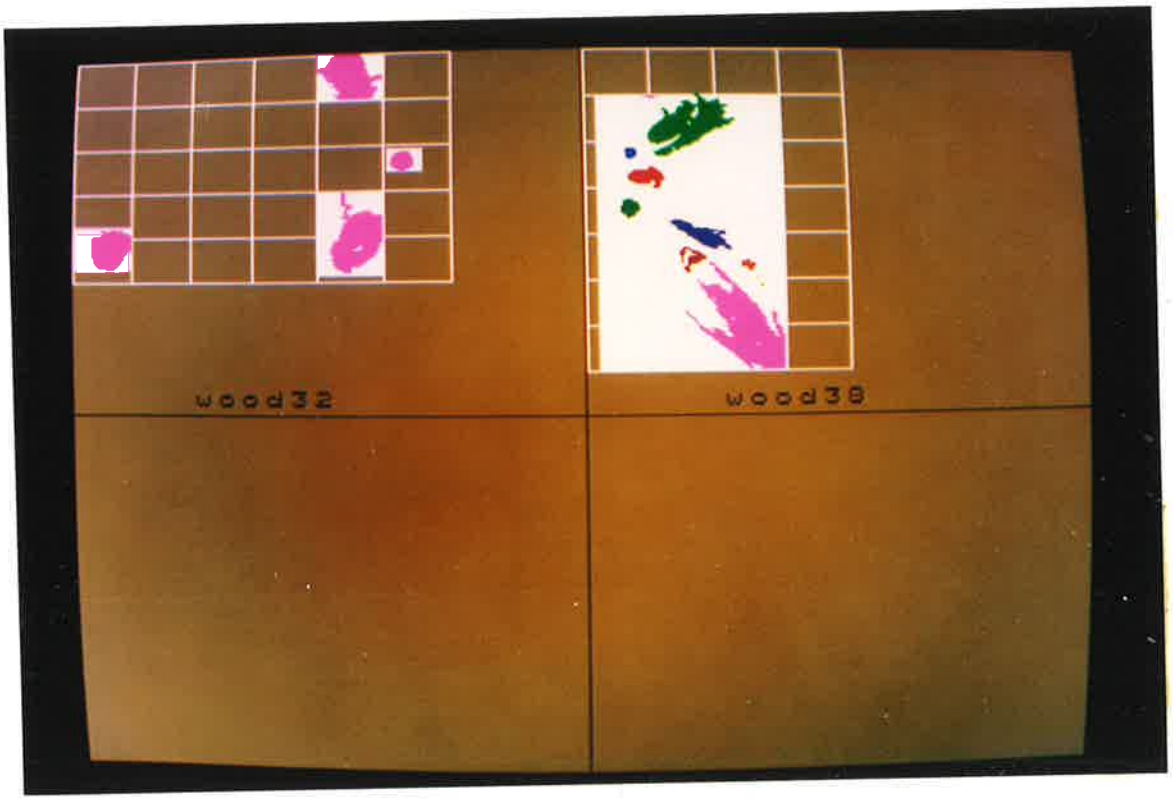

Figure A3A.7 Training set - WOOD32, WOOD38.

BLOB DESCRIPTION - TRAINING SET

LINE STATISTICS METHOD

Presented below is a list of all the blobs identified in the previous figures. The measures are:

area in mm^2, of the blob.

x,y in mm, of the top left corner of a box that encloses the blob. The origin is the top left corner of each image.

dx,dy in mm, the size of the blob. The significant dimension for the grading rules is dy which is the size of the feature.

len/wid the ratio of the length to the width (dx/dy) is a measure of the shape of the blob.

propn the proportion of the enclosing box of the blob that contains blob pixels.

p^2/a ratio of the square of the perimeter of the blob to the area. This is a measure of the shape.

coldiff the difference in average grey levels between the blob pixels and the background pixels in the box that encloses the blob.

The blobs that define features are marked with the following code:

Large knot(hole)	lkn(lho)
Medium knot(hole)	mkn(mho)
Small knot(hole)	skn(sho)
Pin knot(hole)	pkn(pho)
Needle trace	ntr
Pith	pth
Resin	rsn
Bark	brk
Growth ring	grn

****

wood01

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	88.4	35.2	62.0	10.1	11.7	0.869	0.749	22.7	118	sho
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	263.3	156.7	92.1	17.1	32.0	0.534	0.483	46.8	77	lkn
2	12.4	148.7	104.9	10.7	3.8	2.835	0.310	58.1	28	

****

wood02

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	47.9	31.4	96.6	8.5	7.9	1.080	0.711	22.9	81	pkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	582.8	106.1	33.5	28.8	28.6	1.007	0.709	24.0	103	mho

****

wood03

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1021.3	108.7	26.3	55.4	30.8	1.798	0.598	105.7	36	lkn
2	13.4	140.7	52.3	6.4	4.1	1.546	0.508	46.8	11	
3	13.8	133.8	59.8	5.9	6.8	0.866	0.348	48.8	18	

****

wood04

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1365.6	90.1	80.8	67.7	36.1	1.875	0.559	51.1	67	lkn
2	21.6	153.5	103.8	9.6	13.5	0.709	0.167	112.0	7	

****

wood05

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	21.4	231.3	7.9	7.5	4.9	1.527	0.588	27.3	20	
2	2686.3	0.0	16.5	238.8	16.9	14.113	0.665	87.7	75	pth
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	68.3	0.0	138.0	18.1	6.4	2.835	0.590	71.4	26	

****

wood06

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	3854.6	68.2	0.0	123.1	144.4	0.853	0.217	463.6	87	lkn
2	10.0	68.2	22.2	3.7	5.3	0.709	0.510	35.3	43	
3	13.2	121.0	56.4	21.9	1.1	19.373	0.537	117.3	35	
4	22.0	81.5	71.1	5.3	13.2	0.405	0.314	98.3	26	
5	87.6	123.1	105.3	22.4	33.8	0.662	0.116	169.3	29	
6	26.5	173.8	118.1	30.9	2.3	13.703	0.379	132.0	33	
7	25.5	188.1	129.7	16.5	2.6	6.278	0.585	45.5	34	
8	28.3	179.6	131.6	14.4	8.3	1.740	0.237	92.2	31	

****

wood07

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	140.5	53.8	29.0	106.6	17.3	6.163	0.076	345.3	96	rsn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	27.3	150.8	0.0	16.5	4.9	3.380	0.337	56.9	41	rsn

****

wood08

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	57.7	0.0	26.7	9.6	8.6	1.109	0.696	25.7	51	pkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	904.6	97.5	20.7	69.8	31.2	2.237	0.415	120.7	59	lkn
2	26.9	82.1	24.1	4.8	7.9	0.608	0.709	28.7	47	pkn
3	22.4	90.1	53.4	6.4	4.5	1.418	0.778	22.3	51	pkn
4	15.8	116.7	68.4	6.9	4.5	1.536	0.506	31.6	42	
5	92.2	109.8	71.4	17.6	13.9	1.264	0.377	51.6	27	lkn
6	50.3	124.7	74.8	13.3	8.6	1.541	0.437	46.5	36	
7	16.2	141.8	89.9	3.2	7.5	0.425	0.675	38.7	15	
8	48.9	158.8	103.0	17.6	8.6	2.034	0.321	71.4	51	pkn
9	23.8	164.7	112.0	22.9	7.9	2.903	0.132	137.7	50	

****

wood09

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	682.0	30.9	26.7	207.9	9.4	22.114	0.349	245.5	73	brk
2	373.6	164.2	41.4	74.6	12.4	6.014	0.403	168.2	35	
3	4333.4	0.0	54.9	238.8	57.5	4.151	0.315	217.0	107	brk
4	27.1	97.0	77.1	18.1	3.0	6.025	0.496	57.4	22	
5	12.0	18.7	97.8	4.8	3.4	1.418	0.741	21.6	34	

****

wood10

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	81.4	174.3	120.3	44.2	11.7	3.795	0.158	207.1	34	
2	167.5	196.7	121.4	42.1	15.8	2.666	0.252	139.9	29	
3	1164.6	54.9	135.4	183.9	33.1	5.557	0.191	184.0	114	brk
4	17.4	29.3	144.0	20.3	2.3	8.978	0.382	97.3	40	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	2735.2	73.6	53.8	90.6	57.2	1.585	0.528	71.5	53	lkn
2	98.2	174.3	60.9	15.5	15.8	0.979	0.402	49.7	34	grn
3	205.6	167.4	78.2	22.9	32.7	0.701	0.274	111.3	23	grn
4	12.4	188.7	107.9	7.5	3.0	2.481	0.554	31.2	18	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	68.5	208.9	88.7	29.8	12.4	2.406	0.185	96.9	39	grn

****

wood11

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	131.5	108.7	126.3	13.9	13.2	1.053	0.721	22.7	63	pkn
2	344.7	127.9	153.8	110.9	11.7	9.511	0.267	226.4	78	brk
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	81.2	119.4	74.4	11.7	9.0	1.299	0.767	22.8	38	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	211.8	155.6	20.7	19.7	17.7	1.116	0.608	28.0	63	skn

****

wood12

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	66.7	53.8	104.9	9.6	9.0	1.063	0.771	21.2	85	pho
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	64.9	117.3	51.5	22.9	5.3	4.354	0.538	41.5	119	brk
2	10.8	99.1	56.8	11.7	1.9	6.237	0.491	54.0	12	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	48.9	177.0	29.3	10.1	10.5	0.962	0.459	59.0	34	
2	27.1	189.2	41.0	6.4	7.1	0.895	0.592	36.3	38	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1644.3	170.6	120.3	68.2	48.1	1.418	0.501	67.5	61	lkn
2	12.8	229.2	140.2	9.6	2.6	3.645	0.508	39.1	25	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	352.1	204.7	0.0	25.1	24.1	1.041	0.584	64.3	29	lkn

****

wood13

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	173.0	40.5	53.8	16.5	13.5	1.221	0.773	34.3	40	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	

1	339.1	125.8	126.3	49.6	35.3	1.402	0.194	122.9	98	lkn

****

wood14

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	208.0	58.6	136.9	20.3	16.2	1.253	0.635	25.3	51	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	11.0	102.9	100.8	4.3	3.8	1.134	0.688	26.3	23	
2	281.2	71.4	103.0	29.3	15.8	1.856	0.607	34.5	43	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	11.2	165.8	3.8	10.1	1.9	5.387	0.589	41.1	26	
2	28.1	178.0	4.5	7.5	8.6	0.863	0.435	57.9	24	
3	11.6	188.7	7.9	3.7	4.1	0.902	0.753	22.3	22	
4	15.8	175.9	12.4	6.9	6.4	1.084	0.357	65.6	16	
5	223.1	154.6	18.8	29.3	29.3	1.000	0.259	111.3	55	lkn
6	17.6	145.5	35.0	6.4	5.6	1.134	0.489	38.2	40	

****

wood15

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	211.0	72.5	105.7	19.2	14.7	1.309	0.750	21.9	36	skn

****

wood16

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	92.4	89.0	30.8	11.2	19.6	0.572	0.422	114.8	41	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	536.7	107.1	120.7	28.8	23.7	1.215	0.787	21.9	41	mkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	156.9	127.9	123.3	22.4	13.2	1.701	0.533	229.6	18	

****

wood17

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	70.1	62.4	57.9	12.3	9.8	1.254	0.585	39.8	22	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	44.3	72.0	139.1	16.0	5.3	3.038	0.526	56.8	30	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	116.4	139.1	75.2	14.4	14.7	0.981	0.552	54.5	44	mkn

****

wood18

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	2616.9	75.2	0.0	140.7	71.1	1.980	0.262	262.1	113	lkn
2	30.9	119.4	28.2	13.9	4.9	2.835	0.456	39.5	62	

****

wood19

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	460.9	16.0	96.3	33.6	21.1	1.595	0.652	37.1	40	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	491.6	48.0	120.7	34.6	21.1	1.645	0.674	35.2	33	mkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	51.7	85.8	0.0	19.7	16.9	1.166	0.155	106.8	53	

2	246.1	101.3	0.0	34.1	30.1	1.134	0.240	125.1	28	lkn
3	146.1	102.9	11.3	30.9	26.7	1.158	0.177	170.0	30	
4	19.4	99.7	16.5	6.4	6.0	1.063	0.505	32.3	38	

****

wood21

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	28.9	30.4	25.6	8.0	5.3	1.519	0.686	23.4	30	pkn
2	90.2	41.6	29.3	14.9	11.7	1.280	0.518	37.6	52	skn
3	14.4	37.8	45.9	4.8	4.1	1.160	0.727	22.2	64	pkn
4	13.0	54.9	62.4	6.4	4.9	1.309	0.417	44.9	22	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	153.9	50.1	96.3	11.7	21.4	0.547	0.612	39.4	38	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	61.3	71.4	89.5	14.9	18.0	0.827	0.228	75.5	60	
2	32.9	69.8	115.8	13.3	13.9	0.958	0.177	93.8	19	
3	32.5	117.8	125.6	6.9	7.1	0.970	0.656	30.2	16	
4	879.0	82.6	130.1	60.8	38.4	1.584	0.377	92.2	56	lkn
5	32.7	136.4	152.7	8.5	9.0	0.945	0.424	88.3	31	

****

wood22

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	30.7	103.4	24.1	8.0	8.6	0.924	0.443	37.8	56	
2	2696.7	0.0	30.5	238.8	28.6	8.356	0.395	114.6	76	pth
3	11.0	28.8	58.3	3.2	4.9	0.654	0.705	26.3	41	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	20.4	97.0	85.7	5.3	4.9	1.090	0.785	22.6	26	

****

wood24

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	298.8	58.6	83.5	25.6	17.7	1.448	0.661	36.1	32	skn
2	598.2	70.4	104.5	32.0	26.7	1.198	0.701	39.6	36	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	239.3	110.9	18.4	25.6	21.4	1.194	0.436	60.2	51	lkn
2	17.0	98.1	25.6	9.1	5.6	1.607	0.333	48.2	38	
3	149.7	98.6	38.4	22.9	25.6	0.896	0.255	78.4	47	

****

wood26

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	112.8	85.3	57.2	14.9	11.7	1.280	0.649	25.6	106	sho
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	209.6	106.6	24.1	18.1	14.3	1.268	0.810	31.0	32	skn

****

wood31

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	18.0	14.9	136.9	4.8	5.3	0.911	0.714	25.6	48	pkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	122.6	42.1	12.8	12.8	15.0	0.851	0.638	39.8	55	lkn
2	11.2	42.1	37.6	3.2	5.6	0.567	0.622	31.5	21	
3	39.3	34.1	41.0	6.4	12.8	0.500	0.480	45.1	40	

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	14.2	45.8	111.7	4.3	4.1	1.031	0.807	20.3	36	pkn

****

wood32

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	411.2	10.1	90.6	24.0	22.2	1.081	0.773	23.6	57	mkn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	530.5	136.4	0.0	34.1	24.1	1.418	0.646	33.5	54	mkn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	662.9	142.8	72.6	27.7	43.2	0.641	0.553	48.4	48	mkn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	101.6	173.8	51.5	12.8	10.5	1.215	0.754	21.3	48	skn

****

wood38

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	11.4	28.8	24.1	10.1	1.9	5.387	0.600	40.4	9	
2	742.3	32.5	24.1	46.9	32.3	1.451	0.489	79.3	54	skn
3	34.5	20.3	51.5	6.9	6.0	1.152	0.827	19.6	59	pkn
4	134.3	21.9	61.3	19.7	11.7	1.692	0.584	36.3	55	pkn
5	79.6	18.1	78.2	10.7	10.2	1.050	0.735	23.2	63	pkn
6	217.6	42.6	89.1	33.0	16.9	1.953	0.389	56.6	36	
7	78.0	48.0	103.8	14.4	13.2	1.094	0.412	107.0	16	
8	26.9	79.4	110.2	8.0	6.0	1.329	0.558	36.6	20	
9	1215.5	56.5	112.0	45.8	56.4	0.813	0.470	115.8	37	lkn
10	10.0	86.3	119.2	5.3	5.6	0.945	0.333	58.3	14	
11	11.0	49.6	137.2	7.5	6.4	1.167	0.231	69.9	26	

FEATURE DESCRIPTION - TRAINING SET

LINE STATISTICS METHOD

This list contains all the features listed in appendix 1A and describes how well they are extracted using the line statistics method to define the feature area extent. The features are listed in the catagories of table 9.2 and list DY, the feature size as determined by a Quality Control Officer, and DY', the size of the blob as determined by the discrimination process (both im millimetres).

TRAINING SET

LARGE KNOTS

IMAGE	DY	DY'	COMMENT
WOOD01	37.2	32.0	bark only, 1 extra small blob
WOOD03	33.5	30.8	2 extra small blobs, well defined
WOOD04	33.1	36.1	well defined, one extra blob
WOOD06	144.4	144.4	7 extra blobs, spike knot, well defined
WOOD08	30.1	31.2	4 g-ring blobs, size well defined
WOOD08	42.1	13.9	bark only, poorly defined
WOOD10	44.0	57.2	well defined but extended by g-ring slightly
WOOD12	55.3	24.1	2 bark blobs, feature undefined
WOOD12	30.8	48.1	1 extra blob, extended by g-ring
WOOD13	32.7	35.3	bark only, defines feature size
WOOD14	47.8	29.3	6 bark blobs, feature poorly defined
WOOD15	37.2		classed clear
WOOD18	78.2	71.1	bark only, 1 extra blob, extent defined well
WOOD19	41.0	30.1	4 bark blobs, feature poorly defined
WOOD21	107.2	38.4	5 blobs, feature poorly defined
WOOD24	47.4	21.4	3 bark blobs, feature poorly defined
WOOD31	52.3	15.0	bark only, size poorly defined
WOOD38	77.8	56.4	6 blobs, passably defined

MEDIUM KNOTS

IMAGE	DY	DY'	COMMENT
WOOD02	28.6	28.6	hole, well defined
WOOD13	25.6		not detected, classed clear
WOOD15	25.2		classed clear
WOOD16	24.8	23.7	defined in two areas, size well defined
WOOD17	27.1	14.7	bark only, does not define feature size
WOOD19	22.9	21.1	defined well
WOOD32	20.7	22.2	defined well
WOOD32	21.4	24.1	defined well
WOOD32	27.1	43.2	defined well but g-ring extends feature size

SMALL KNOTS

IMAGE	DY	DY'	COMMENT - defined well unless otherwise
WOOD01	11.3	11.7	hole
WOOD11	13.5	9.0	

WOOD11	16.5	17.7	
WOOD13	13.9	13.5	
WOOD14	14.7	16.2	
WOOD14	15.4	15.8	
WOOD15	12.4		classed clear
WOOD15	14.3	14.7	
WOOD16	17.3	19.6	bark only, extent defined well
WOOD17	18.0	9.8	size defined porrly
WOOD17	13.2	5.3	defined under size
WOOD19	19.6	21.1	
WOOD21	12.4	11.7	
WOOD24	13.5	17.7	
WOOD24	19.9	26.7	
WOOD26	11.3	11.7	hole
WOOD26	16.5	14.3	
WOOD32	10.5	10.5	
WOOD38	16.5	32.3	

PIN KNOTS

IMAGE	DY	DY'	COMMENT - well defined unless otherwise
WOOD02	6.0	7.9	
WOOD08	7.9	8.6	
WOOD08	7.5	7.9	
WOOD08	5.6	4.5	
WOOD08	6.4	8.6	
WOOD11	9.4	13.2	
WOOD12	8.6	9.0	hole
WOOD14	7.1		classed clear
WOOD21	5.3	5.3	
WOOD21	4.1	4.1	
WOOD31	5.6	5.3	
WOOD31	4.9	4.1	
WOOD38	6.0	6.0	
WOOD38	8.3	11.7	
WOOD38	9.4	10.2	

PITH

IMAGE	COMMENT
WOOD05	well defined
WOOD22	well defined

RESIN AND BARK

IMAGE	COMMENT
WOOD07	resin, well defined
WOOD07	resin, well defined
WOOD09	bark, well defined
WOOD09	bark, well defined but extended by false positive area
WOOD10	bark, well defined, 2 extra g-ring blobs
WOOD11	bark, extended by g-rings
WOOD12	bark, small piece defined well

APPENDIX 3B - FEATURE DISCRIMINATION

TEST SET - LINE STATISTICS METHOD

The photographs presented in this appendix are the processed results of the test set of the WOOD library of images. For the purposes of display the original images are reduced to one-quarter size by displaying only every second pixel in every second row of each image. In this way four images are displayed on one screen and photographed.

Each figure is the result of processing using the line statistics method to determine the feature area extent. Each feature area is processed using the modification of the adaptive threshold method for feature areas. The grid lines are the boundaries of the local areas of size 64 * 64 pixels. The white rectangles are the feature areas and each coloured blob is separate from every other. Vacant areas that can be inferred by the absence of grid lines are feature areas that have been detected and classed as clear.

Following the figures is a list of the location and size of every blob with various descriptive shape measures. This is followed by a description of how well each feature is described.

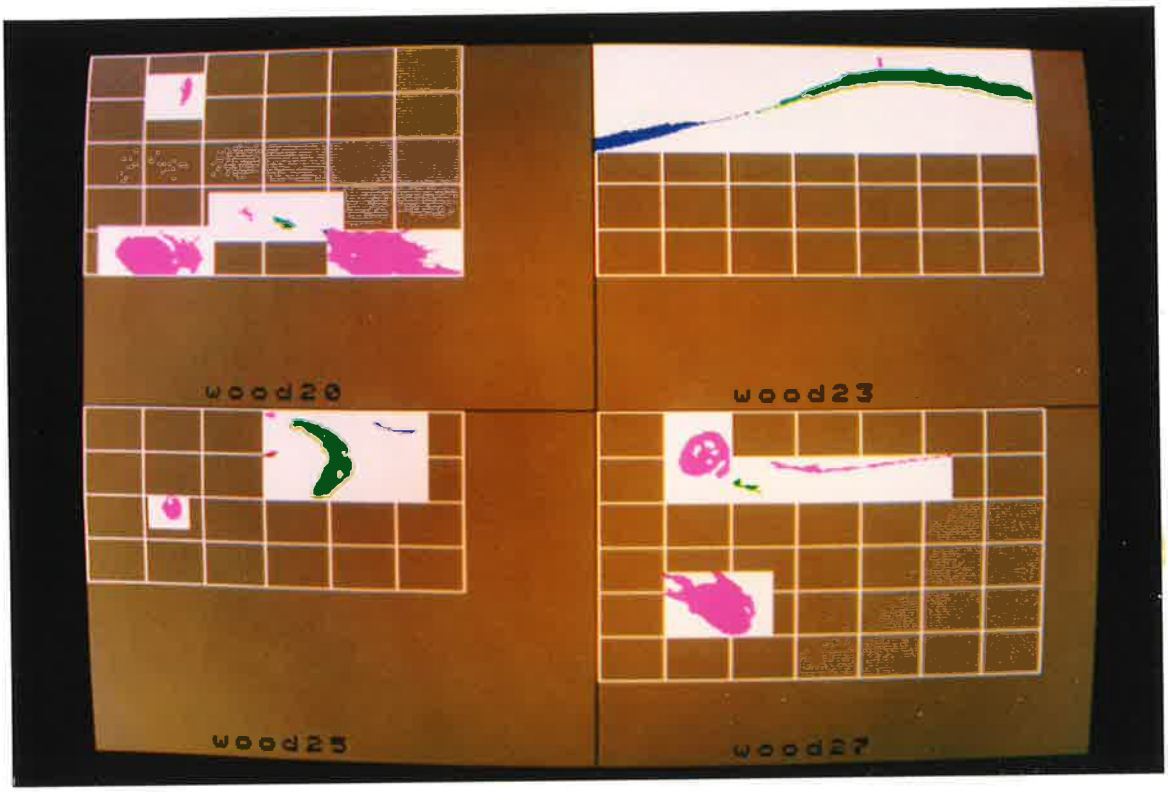

Figure A3B.1 Test set - WOOD20, WOOD23, WOOD25, WOOD27.

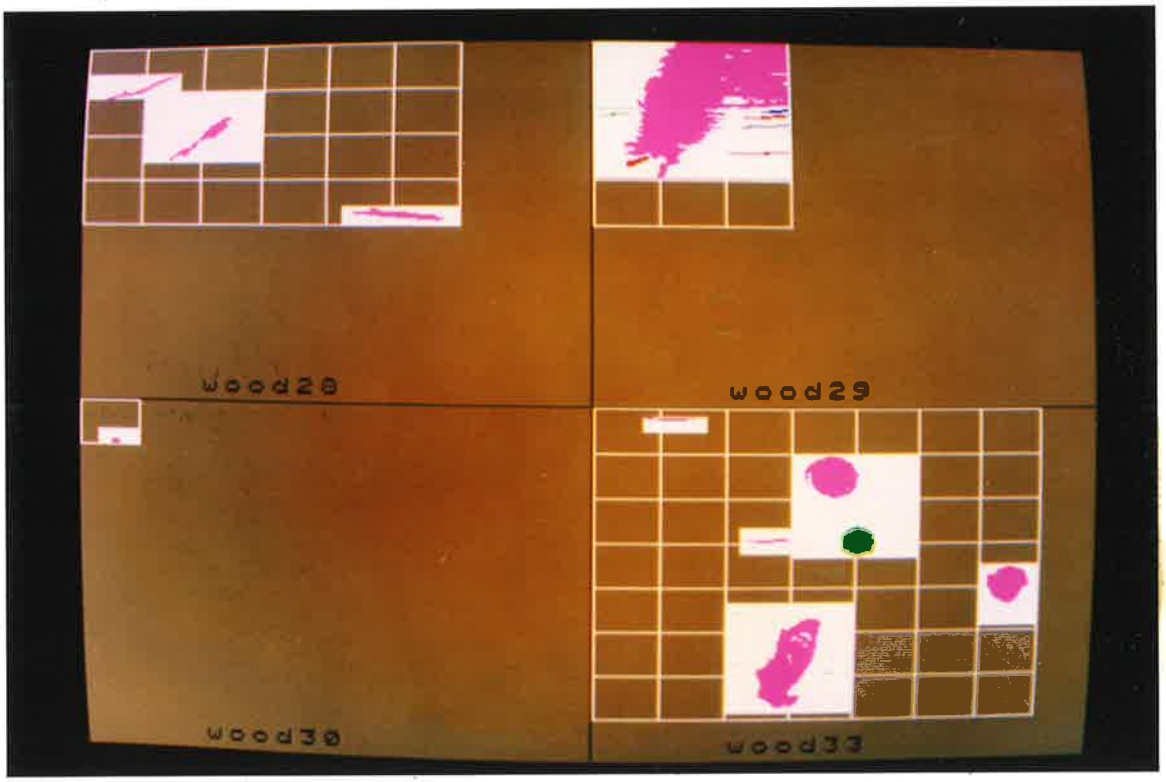

Figure A3B.2 Test set - WOOD28, WOOD29, WOOD30, WOOD33.

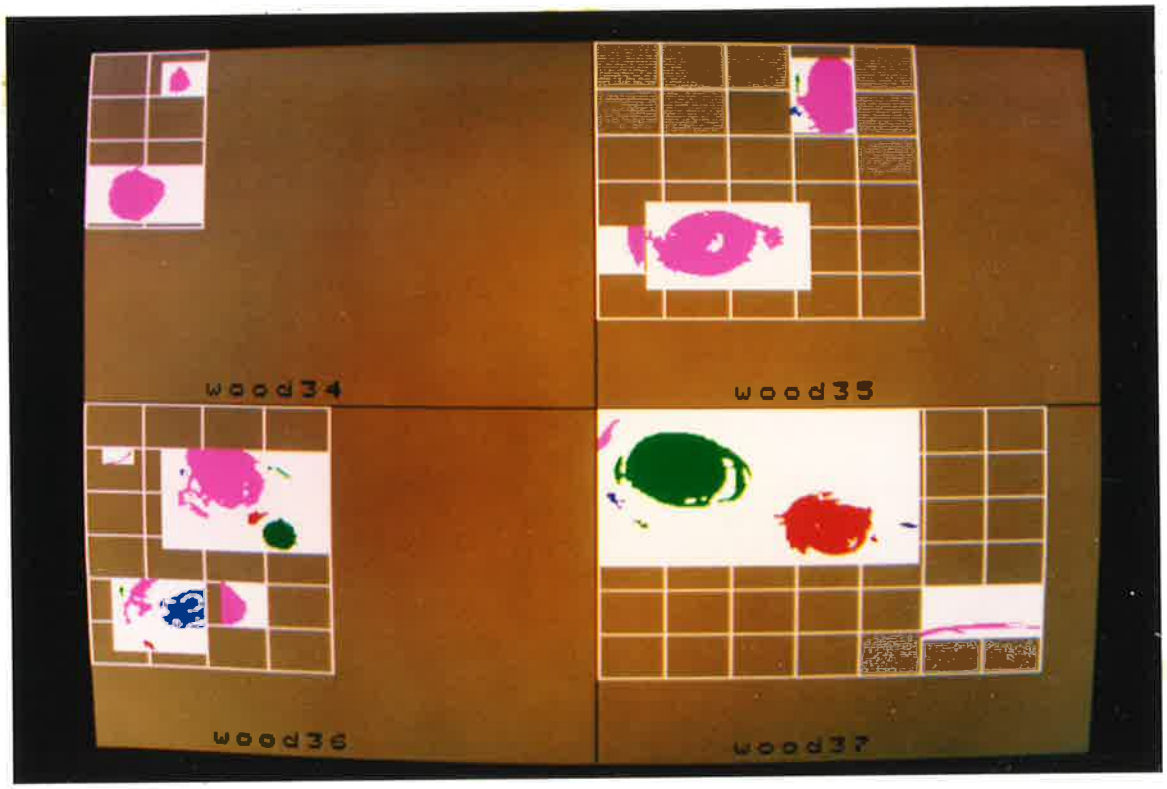

Figure A3B.3 Test set - WOOD34, WOOD35, WOOD36, WOOD37.

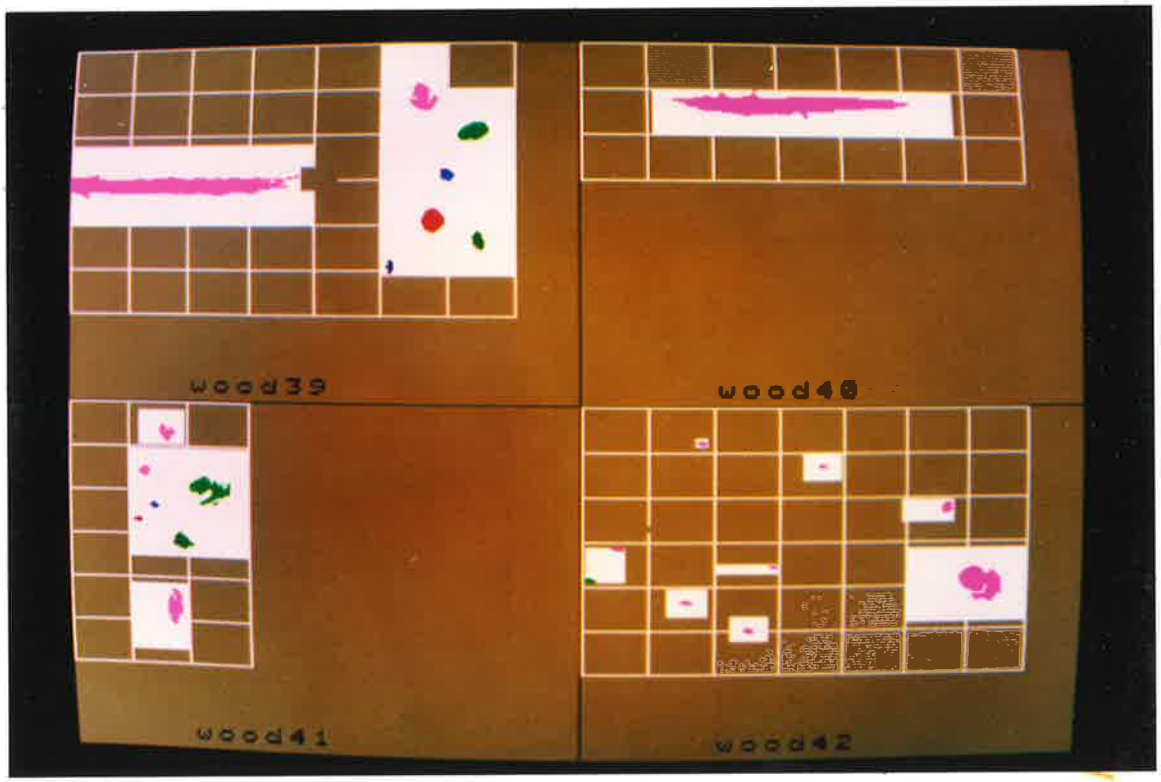

Figure A3B.4 Test set - WOOD39, WOOD40, WOOD41, WOOD42.

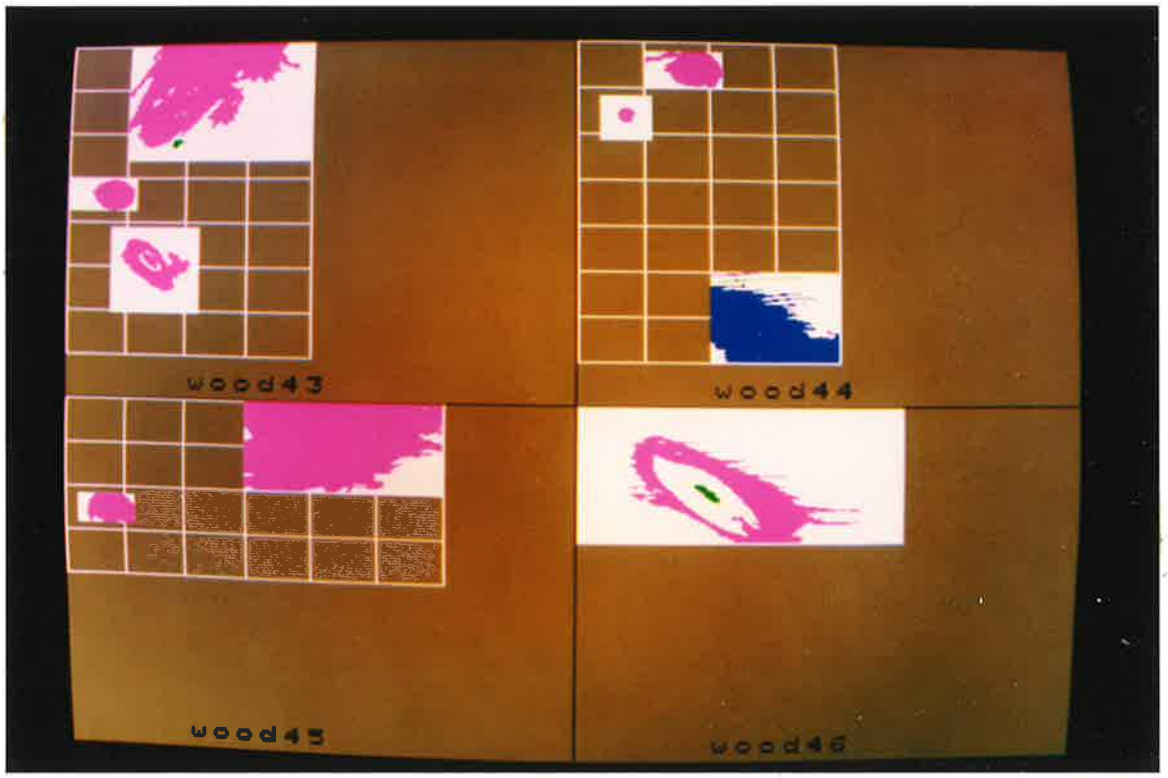

Figure A3B.5 Test set - WOOD43, WOOD44, WOOD45, WOOD46.

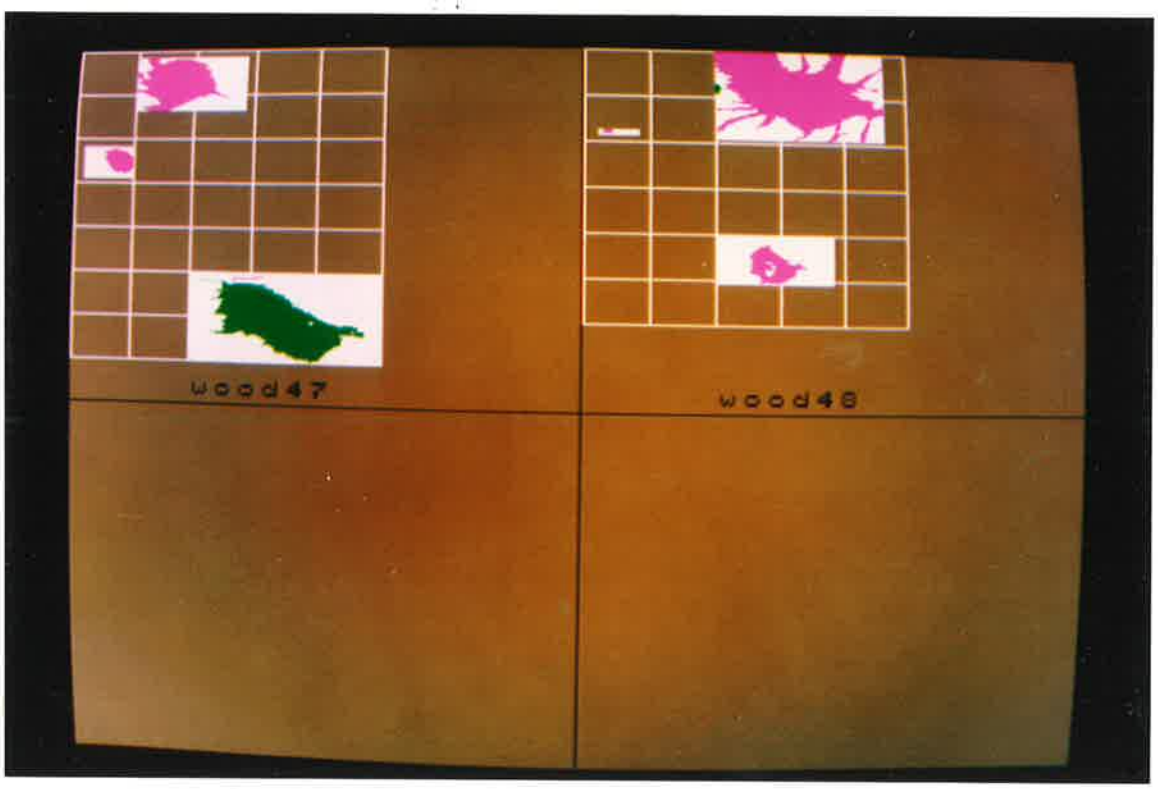

Figure A3B.6 Test set - WOOD47, WOOD48.

BLOB DESCRIPTION - TEST SET

LINE STATISTICS METHOD

Presented below is a list of all the blobs identified in the previous figures. The measures are:

area in mm^2, of the blob.

x,y in mm, of the top left corner of a box that encloses the blob. The origin is the top left corner of each image.

dx,dy in mm, the size of the blob. The significant dimension for the grading rules is dy which is the size of the feature.

len/wid the ratio of the length to the width (dx/dy) is a measure of the shape of the blob.

propn the proportion of the enclosing box of the blob that contains blob pixels.

p^2/a ratio of the square of the perimeter of the blob to the area. This is a measure of the shape.

coldiff the difference in average grey levels between the blob pixels and the background pixels in the box that encloses the blob.

The blobs that define features are marked with the following code:

Large knot(hole)	lkn(lho)
Medium knot(hole)	mkn(mho)
Small knot(hole)	skn(sho)
Pin knot(hole)	pkn(pho)
Needle trace	ntr
Pith	pth
Resin	rsn
Bark	brk
Growth ring	grn

****

wood20

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	870.4	17.1	96.3	51.7	24.1	2.148	0.700	53.5	37	mkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	68.3	53.8	12.8	7.5	15.8	0.473	0.580	42.2	51	lkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	23.8	89.0	83.8	8.5	6.8	1.260	0.413	48.5	27	
2	42.5	107.1	88.4	12.8	7.1	1.791	0.465	39.9	21	
3	61.5	132.7	95.1	12.3	7.1	1.716	0.703	27.6	38	grn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1109.7	136.4	96.3	68.2	24.1	2.835	0.676	44.4	80	mkn

****

wood23

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	10.6	150.8	5.6	3.2	5.3	0.608	0.631	30.2	34	
2	1011.3	82.6	10.9	156.2	25.2	6.199	0.257	130.0	76	pth

3	348.9	0.0	35.0	81.5	21.1	3.873	0.203	105.2	79	pth

****

wood25

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	110.6	44.8	48.1	11.7	13.2	0.891	0.717	23.5	83	sho
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	16.4	102.3	1.5	7.5	3.0	2.481	0.732	28.1	8	
2	567.4	115.1	5.3	34.6	42.1	0.823	0.389	59.4	77	mkn
3	41.3	159.4	5.6	22.4	6.0	3.721	0.307	65.3	66	
4	16.2	102.3	22.2	8.5	3.4	2.520	0.563	30.9	48	

****

wood27

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	477.6	41.6	10.2	26.7	25.6	1.042	0.701	38.3	27	mkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	932.1	34.1	85.7	48.0	34.2	1.402	0.568	48.3	53	mkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	175.2	88.5	24.1	97.5	9.4	10.376	0.191	248.5	61	rsn
2	45.7	68.8	35.3	16.0	7.9	2.025	0.362	63.2	28	

****

wood28

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	71.3	10.1	14.7	43.2	12.8	3.377	0.129	159.1	76	rsn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	136.7	49.6	38.4	36.8	22.9	1.603	0.162	110.1	60	lkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	131.3	149.2	87.2	48.5	7.1	6.789	0.379	78.0	80	rsn

****

wood29

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	3269.9	14.9	0.0	87.4	71.4	1.224	0.524	80.5	74	lho
2	19.0	0.0	34.2	21.3	1.1	18.901	0.792	77.9	17	
3	20.6	86.9	34.6	15.5	2.3	6.852	0.592	47.6	13	
4	26.3	77.3	37.2	25.1	2.6	9.518	0.398	89.0	23	
5	12.6	2.1	37.6	16.5	1.9	8.789	0.406	82.3	30	
6	20.4	76.2	42.5	26.1	2.3	11.577	0.347	118.6	22	
7	27.1	70.4	56.8	32.0	1.9	17.011	0.450	129.1	22	
8	29.9	16.5	59.8	11.7	5.6	2.079	0.452	38.8	20	

****

wood30

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	15.0	17.6	20.7	5.3	3.4	1.575	0.833	19.3	36	skn

****

wood33

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	30.7	25.6	4.5	27.2	1.9	14.459	0.600	91.0	29	rsn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	40.7	78.4	68.4	31.4	3.4	9.293	0.382	91.1	45	rsn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	923.7	84.7	113.2	34.6	48.9	0.709	0.545	40.9	52	lkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	473.6	109.3	25.2	28.8	22.6	1.276	0.729	31.3	40	mkn
2	226.5	128.5	63.2	18.7	15.8	1.181	0.769	22.7	35	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	339.5	210.0	85.7	24.0	20.3	1.181	0.697	23.1	56	skn

****

wood34

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	710.8	12.3	62.0	33.6	30.5	1.103	0.695	25.7	71	mkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	129.7	46.4	7.9	13.3	14.3	0.933	0.681	24.5	62	skn

****

wood35

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	152.3	15.5	96.3	9.1	21.1	0.430	0.798	30.4	41	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1514.7	25.6	87.6	72.0	34.6	2.080	0.609	47.6	37	lkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	937.7	105.5	7.5	30.9	38.7	0.798	0.783	27.7	54	mkn
2	20.6	104.5	14.7	3.2	11.3	0.284	0.572	65.3	9	
3	19.0	102.3	32.7	6.4	5.6	1.134	0.528	43.1	20	

****

wood36

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	19.4	12.8	25.6	14.4	7.1	2.014	0.189	87.3	70	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	209.0	19.7	96.3	21.3	25.2	0.846	0.389	94.5	32	grn
2	14.8	17.6	99.6	3.2	7.5	0.425	0.617	45.5	14	
3	409.4	40.5	100.8	27.7	22.2	1.249	0.666	54.6	25	mkn
4	20.4	30.4	130.5	8.5	5.6	1.512	0.425	48.0	51	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1194.8	53.3	24.1	49.6	37.2	1.332	0.648	67.4	31	mkn
2	24.9	103.9	32.3	13.9	6.4	2.168	0.281	65.3	24	
3	11.0	54.9	35.0	5.3	4.1	1.289	0.500	49.2	7	
4	45.9	92.7	56.4	12.3	7.5	1.630	0.498	40.2	33	
5	234.5	100.7	60.2	19.2	17.3	1.109	0.707	25.3	46	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	256.7	77.3	96.3	16.0	24.1	0.664	0.667	34.4	29	

****

wood37

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	73.1	0.0	4.9	11.7	17.3	0.678	0.361	53.7	25	
2	1923.3	7.5	11.7	71.4	40.2	1.775	0.669	61.8	43	lkn
3	39.1	3.2	43.6	11.7	9.8	1.199	0.341	64.3	23	
4	1180.8	90.1	44.0	52.8	34.2	1.542	0.654	45.5	51	mkn
5	19.2	17.6	58.3	8.5	4.9	1.745	0.462	45.4	25	
6	20.0	157.8	61.7	10.1	2.6	3.848	0.752	27.0	29	
7	12.0	141.2	66.6	4.8	4.9	0.981	0.513	35.3	26	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	

1	174.4	170.6	116.9	68.2	7.1	9.550	0.358	152.3	50	rsn

****

wood39

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1001.8	0.0	68.1	129.5	13.2	9.842	0.588	79.9	72	pth
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	129.5	185.0	21.1	16.0	14.7	1.090	0.552	43.7	45	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	114.2	185.0	24.1	16.0	11.7	1.372	0.613	50.1	45	
2	121.4	209.5	41.7	17.1	10.2	1.680	0.701	23.8	46	skn
3	40.7	200.4	66.6	8.0	7.1	1.119	0.712	22.8	51	pkn
4	112.2	191.9	87.2	11.7	13.2	0.891	0.727	23.2	61	skn
5	49.5	216.4	99.3	6.9	10.2	0.683	0.704	27.2	33	skn
6	22.8	173.8	114.7	4.3	7.9	0.540	0.679	31.6	61	pkn

****

wood40

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	906.4	47.4	24.1	128.5	15.4	8.332	0.458	81.2	69	pth

****

wood41

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	54.1	52.2	11.3	9.6	10.2	0.945	0.556	44.8	22	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	28.3	40.0	34.6	6.9	6.0	1.152	0.678	23.9	50	pkn
2	205.2	68.8	40.6	24.5	16.2	1.516	0.518	72.3	37	skn
3	15.2	45.3	54.9	6.9	3.4	2.048	0.650	27.8	27	pkn
4	11.6	36.2	62.8	6.4	3.0	2.126	0.604	27.6	36	pkn
5	74.8	59.7	71.4	12.8	9.4	1.361	0.622	26.8	24	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	127.1	55.4	102.6	10.1	20.3	0.499	0.618	54.6	40	mkn

****

wood42

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	14.6	14.9	74.8	6.9	2.6	2.633	0.802	21.9	37	ntr
2	15.4	1.6	91.0	6.4	3.0	2.126	0.802	22.9	33	ntr
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	12.8	59.2	17.7	5.3	3.0	1.772	0.800	20.3	57	ntr
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	16.8	49.6	103.4	8.0	3.0	2.658	0.700	25.2	48	ntr
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	10.6	96.5	84.6	5.9	2.6	2.228	0.688	24.5	45	ntr
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	15.8	83.7	118.8	5.9	3.8	1.559	0.718	22.3	47	ntr
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	11.6	122.6	30.1	5.9	3.0	1.949	0.659	24.9	41	ntr
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	31.9	190.3	50.8	6.4	6.4	1.001	0.779	22.6	40	pkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	321.1	200.4	86.5	25.1	19.6	1.281	0.655	50.3	54	skn

****

wood43

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	319.5	13.3	72.9	25.1	16.9	1.481	0.754	21.2	48	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	3018.3	34.1	0.0	95.9	55.3	1.736	0.569	79.2	53	lkn
2	27.3	59.2	51.1	9.6	5.3	1.823	0.540	32.0	12	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	592.0	30.9	105.3	40.0	28.2	1.418	0.525	51.0	31	mkn

****

wood44

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	57.3	20.3	34.6	9.6	8.3	1.160	0.722	22.4	55	pkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	492.2	34.6	4.9	41.0	19.2	2.140	0.625	42.7	57	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	25.9	68.2	120.3	25.6	2.3	11.340	0.448	90.4	38	
2	40.5	68.2	122.6	50.6	3.8	13.467	0.213	226.7	43	
3	1997.5	68.2	127.1	68.2	41.4	1.650	0.708	91.3	83	lkn

****

wood45

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	395.8	12.3	52.3	28.2	16.2	1.747	0.867	23.6	59	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	4174.3	102.3	0.0	102.3	48.1	2.126	0.848	36.3	60	lkn

****

wood46

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	2201.5	25.1	14.7	123.1	57.5	2.140	0.311	151.0	91	lkn
2	65.5	59.7	40.2	14.4	10.5	1.367	0.433	38.4	45	

****

wood47

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	170.1	16.5	53.4	17.6	13.2	1.337	0.735	22.4	50	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	942.1	34.1	3.4	53.8	29.7	1.812	0.589	64.8	65	mkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	15.4	91.7	122.6	17.6	2.3	7.797	0.389	79.0	39	
2	1751.6	69.3	124.1	93.8	44.4	2.114	0.421	81.5	110	lho

****

wood48

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	12.4	8.5	41.7	6.9	2.6	2.633	0.681	25.8	63	pkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	2479.8	68.2	0.0	90.1	46.2	1.948	0.595	180.4	70	mho
2	15.2	68.2	16.9	4.8	5.3	0.911	0.603	30.3	17	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	337.7	85.8	101.5	32.0	20.3	1.575	0.520	40.7	71	skn

FEATURE DESCRIPTION - TEST SET

LINE STATISTICS METHOD

This list contains all the features listed in appendix 1A and describes how well they are extracted using the line statistics method to define the feature area extent. The features are listed in the catagories of table 9.2 and list DY, the feature size as determined by a Quality Control Officer, and DY', the size of the blob as determined by the discrimination process (both im millimetres).

TEST SET

LARGE KNOTS

IMAGE	DY	DY'	COMMENT
WOOD20	34.2	15.8	bark only, does not define feature size
WOOD28	40.2	22.9	bark only, does not define feature size well
WOOD29	51.1	71.4	hole, 8 extra g-ring blobs, size defined
WOOD33	42.1	48.9	defined well
WOOD35	35.0	34.6	defined well
WOOD37	40.2	40.2	some resin noise, defined well
WOOD43	55.3	55.3	size defined well
WOOD44	35.7	41.4	2 g-ring blobs, size defined
WOOD45	32.7	48.1	1 extra blob, extended by g-rings, low thresh
WOOD46	51.1	57.5	bark, plus pith blob, size defined well
WOOD47	39.9	44.4	hole, defined well

MEDIUM KNOTS

IMAGE	DY	DY'	COMMENT
WOOD20	21.4	24.1	defined well
WOOD20	22.9	24.1	size defined well
WOOD25	26.3	42.1	bark only, defines large size, 3 extra blobs
WOOD27	24.1	25.6	well defined
WOOD27	29.0	34.2	well defined, size extended by g-ring
WOOD33	24.1	22.6	well defined
WOOD34	27.8	30.5	defined well
WOOD35	28.6	38.7	defined well, 2 extra blobs
WOOD36	26.7	37.2	defined well, size extended by g-ring
WOOD36	29.7	22.2	defined in 2 areas
WOOD37	26.3	34.2	good, extended by g-ring, 1 extra blobs
WOOD41	21.1	20.3	bark only, size defined well
WOOD43	28.2	28.2	defined well
WOOD47	20.7	29.7	defined well, extended by g-ring
WOOD48	20.3	46.2	hole, extended by g-ring

SMALL KNOTS

IMAGE	DY	DY'	COMMENT - defined well unless otherwise
WOOD20	17.3		classed as clear
WOOD25	13.2	13.2	hole
WOOD30	12.0	3.4	defined under-size

WOOD33	18.0	15.8
WOOD33	19.2	20.3
WOOD34	14.3	14.3
WOOD36	14.7	17.3
WOOD39	15.0	14.7
WOOD39	10.2	10.2
WOOD39	12.8	13.2
WOOD39	11.3	10.2
WOOD41	10.2	10.2
WOOD41	12.8	16.2
WOOD41	13.2	9.4
WOOD42	17.3	19.6
WOOD43	17.3	16.9
WOOD44	16.2	19.2
WOOD45	14.3	16.2
WOOD47	12.8	13.2
WOOD48	15.8	20.3

PIN KNOTS

IMAGE	DY	DY'	COMMENT - defined well unless otherwise
WOOD39	7.5	7.1	
WOOD39	8.3	7.9	
WOOD41	6.8	6.0	
WOOD41	4.5	3.4	
WOOD41	3.8	3.0	
WOOD41	4.1		classed as clear
WOOD42	5.6	2.6	
WOOD44	8.3	8.3	
WOOD48	3.0	2.6	

NEEDLE TRACE

IMAGE	DY	DY'	COMMENT
WOOD42	2.6	3.0	needle trace
WOOD42	4.1	3.0	needle trace
WOOD42	2.3		needle trace
WOOD42	2.3	2.6	needle trace
WOOD42	3.8	6.4	needle trace

PITH

IMAGE	COMMENT
WOOD23	defined well
WOOD23	defined well
WOOD39	defined well
WOOD40	defined well

RESIN AND BARK

IMAGE	COMMENT
WOOD27	resin, defined well
WOOD28	resin, defined well
WOOD28	resin, defined well
WOOD33	resin, defined well
WOOD33	resin, defined well
WOOD33	resin, classed as clear
WOOD37	resin, defined well

APPENDIX 4A - FEATURE DISCRIMINATION

TRAINING SET - ADAPTIVE THRESHOLD METHOD

The photographs presented in this appendix are the processed results of the training set of the WOOD library of images. For the purposes of display the original images are reduced to one-quarter size by displaying only every second pixel in every second row of each image. In this way four images are displayed on one screen and photographed.

Each figure is the result of processing using the adaptive threshold method to determine the feature area extent. Each feature area is processed using the modification of the adaptive threshold method for feature areas. The grid lines are the boundaries of the local areas of size 64 * 64 pixels. The white rectangles are the feature areas and each coloured blob is separate from every other. Vacant areas that can be inferred by the absence of grid lines are feature areas that have been detected and classed as clear.

Following the figures is a list of the location and size of every blob with various descriptive shape measures. This is followed by a description of how well each feature is described.

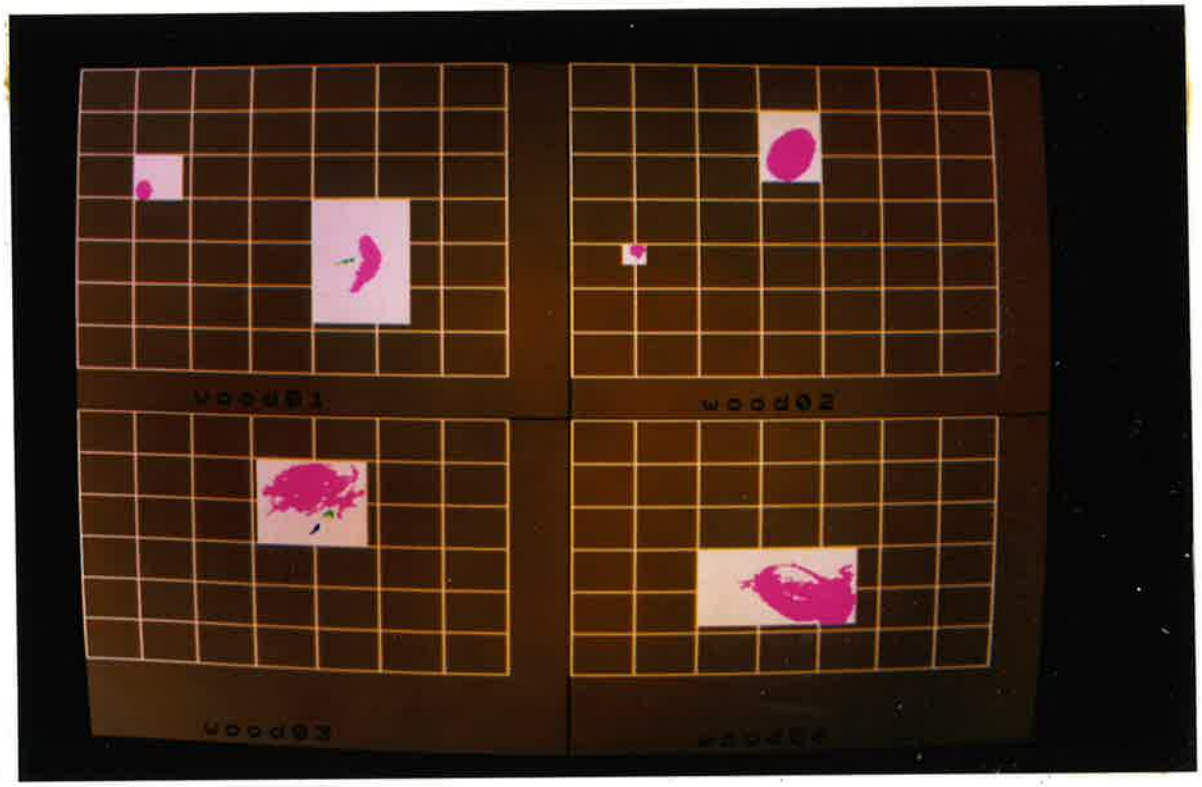

Figure A4A.1 Training set - WOOD01, WOOD02, WOOD03, WOOD04.

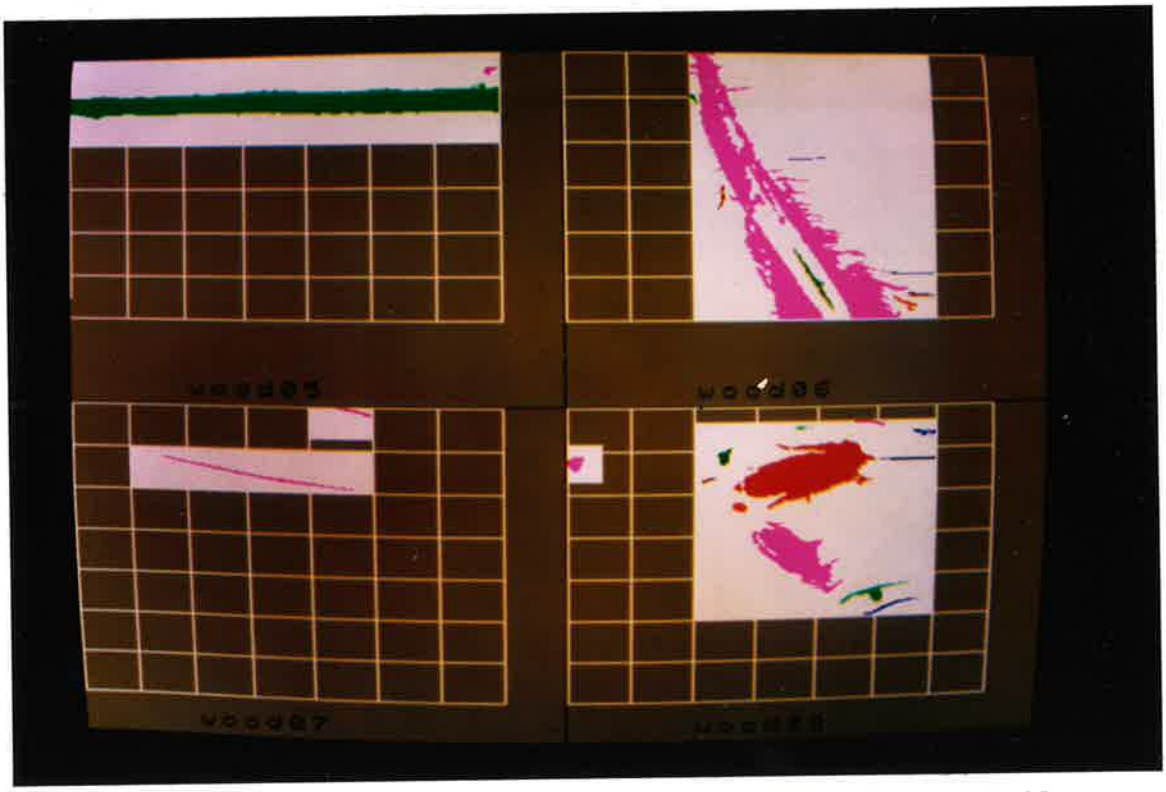

Figure A4A.2 Training set - WOOD05, WOOD06, WOOD07, WOOD08.

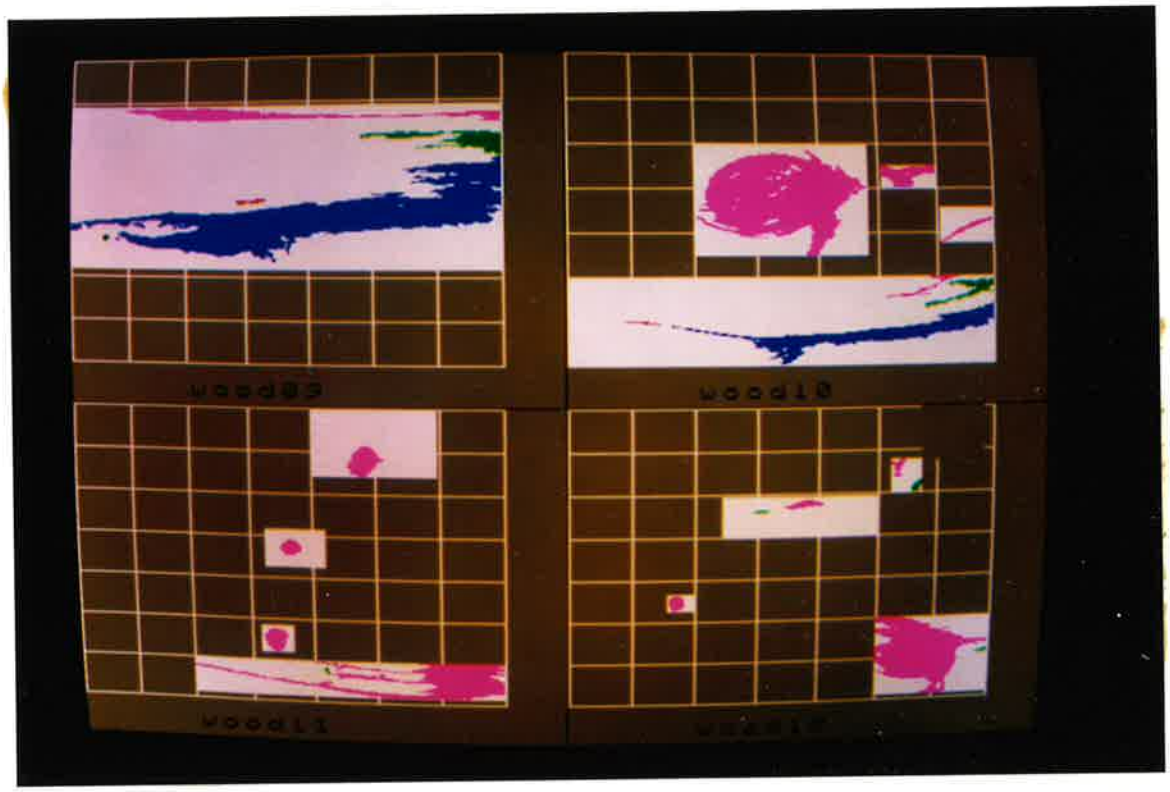

Figure A4A.3 Training set - WOOD09, WOOD10, WOOD11, WOOD12.

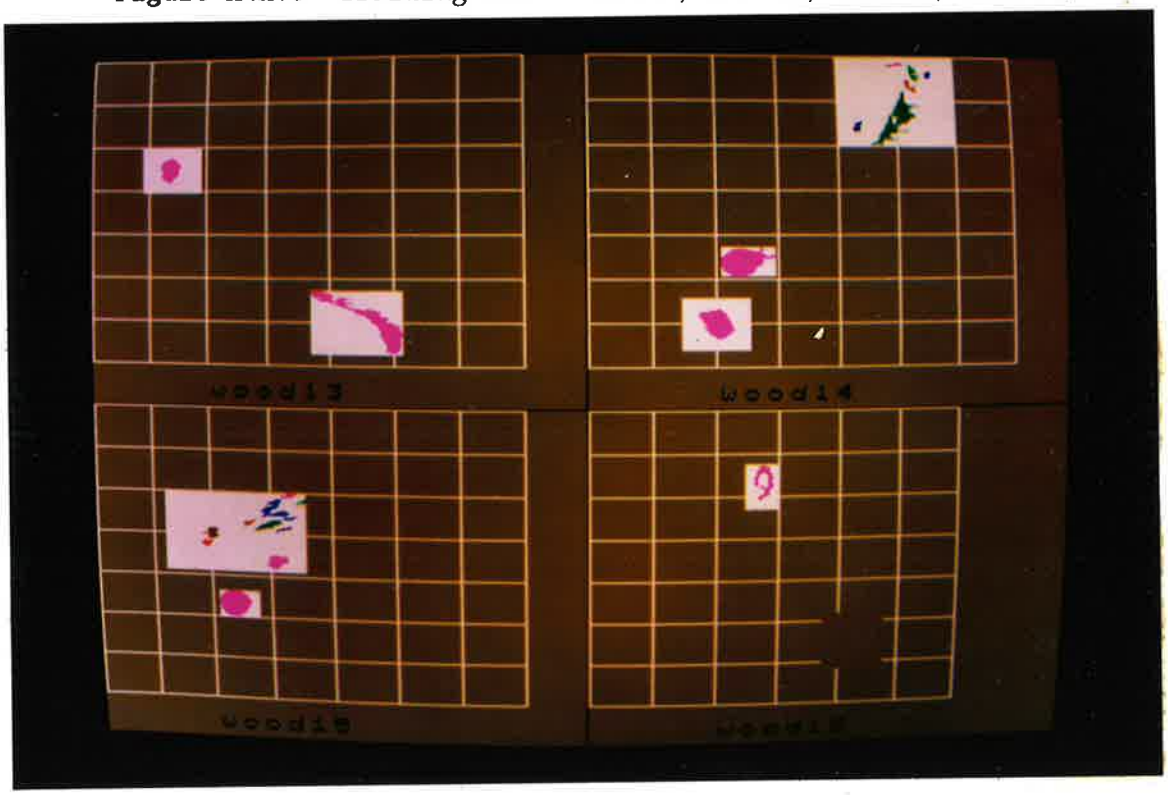

Figure A4A.4 Training set - WOOD13, WOOD14, WOOD15, WOOD16.

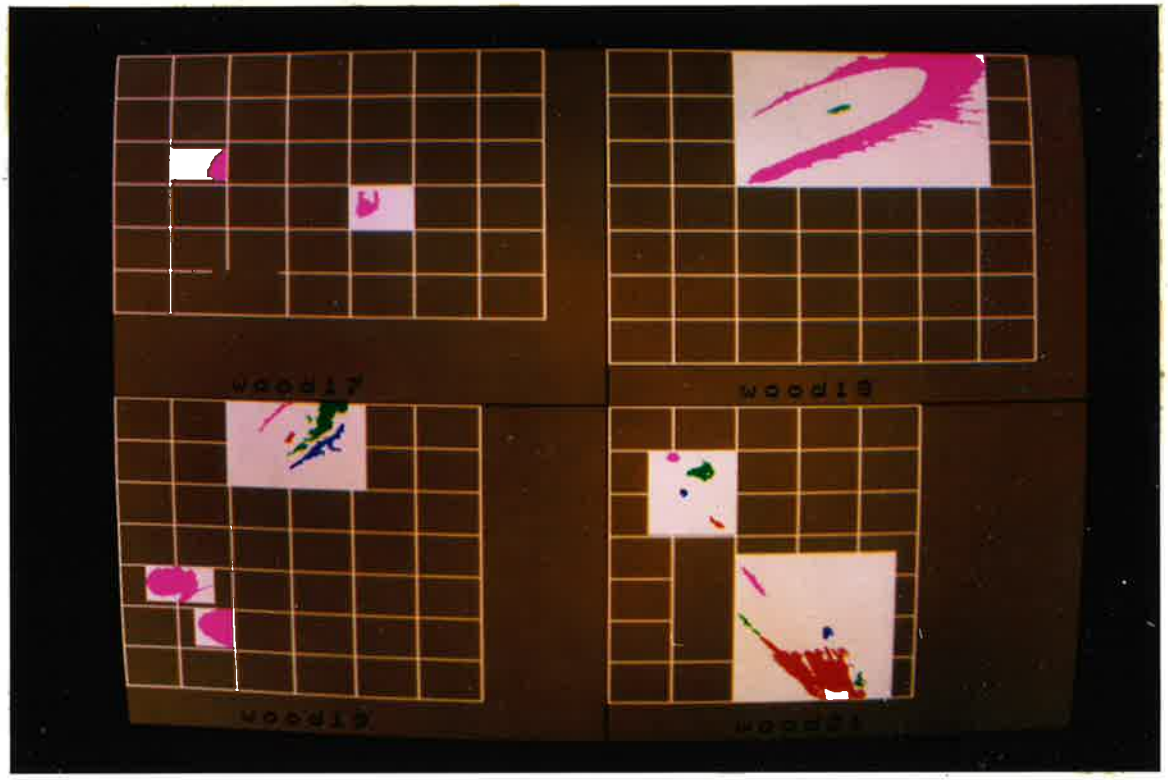

Figure A4A.5 Training set - WOOD17, WOOD18, WOOD19, WOOD21.

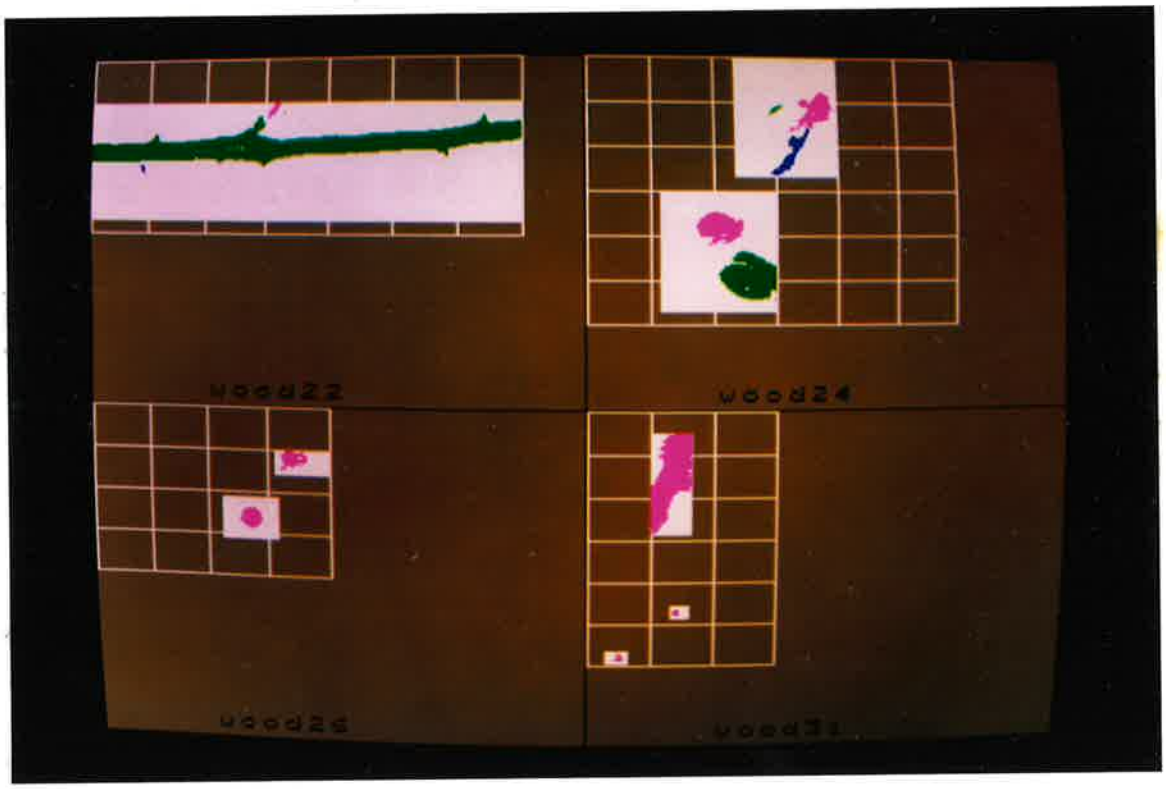

Figure A4A.6 Training set - WOOD22, WOOD24, WOOD26, WOOD31.

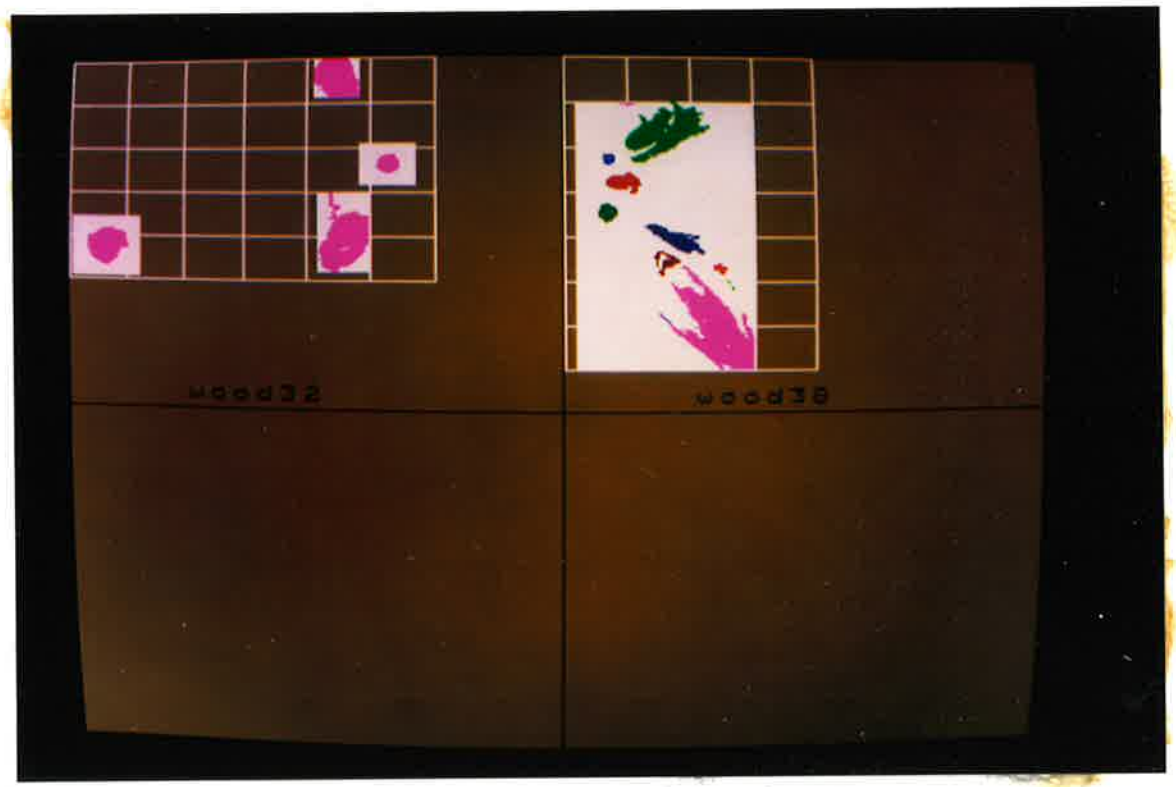

Figure A4A.7 Training set - WOOD32, WOOD38.

BLOB DESCRIPTION - TRAINING SET

ADAPTIVE THRESHOLD METHOD

Presented below is a list of all the blobs identified in the previous figures. The measures are:

area — in mm^2, of the blob.

x,y — in mm, of the top left corner of a box that encloses the blob. The origin is the top left corner of each image.

dx,dy — in mm, the size of the blob. The significant dimension for the grading rules is dy which is the size of the feature.

len/wid — the ratio of the length to the width (dx/dy) is a measure of the shape of the blob.

propn — the proportion of the enclosing box of the blob that contains blob pixels.

p^2/a — ratio of the square of the perimeter of the blob to the area. This is a measure of the shape.

coldiff — the difference in average grey levels between the blob pixels and the background pixels in the box that encloses the blob.

The blobs that define features are marked with the following code:

Large knot(hole)	lkn(lho)
Medium knot(hole)	mkn(mho)
Small knot(hole)	skn(sho)
Pin knot(hole)	pkn(pho)
Needle trace	ntr
Pith	pth
Resin	rsn
Bark	brk
Growth ring	grn

****

wood01

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	88.4	35.2	62.0	10.1	11.7	0.869	0.749	22.7	118	sho
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	263.3	156.7	92.1	17.1	32.0	0.534	0.483	46.8	77	lkn
2	12.4	148.7	104.9	10.7	3.8	2.835	0.310	58.1	28	

****

wood02

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	47.9	31.4	96.6	8.5	7.9	1.080	0.711	22.9	81	pkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	582.8	106.1	33.5	28.8	28.6	1.007	0.709	24.0	103	mho

****

wood03

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1021.3	108.7	26.3	55.4	30.8	1.798	0.598	105.7	36	lkn
2	13.4	140.7	52.3	6.4	4.1	1.546	0.508	46.8	11	
3	13.8	133.8	59.8	5.9	6.8	0.866	0.348	48.8	18	

****

wood04

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1363.4	90.1	80.8	67.2	35.7	1.880	0.568	50.1	67	lkn

****

wood05

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	21.4	231.3	7.9	7.5	4.9	1.527	0.588	27.3	20	
2	2686.3	0.0	16.5	238.8	16.9	14.113	0.665	87.7	75	pth

****

wood06

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	3854.6	68.2	0.0	123.1	144.4	0.853	0.217	463.6	87	lkn
2	10.0	68.2	22.2	3.7	5.3	0.709	0.510	35.3	43	
3	13.2	121.0	56.4	21.9	1.1	19.373	0.537	117.3	35	
4	22.0	81.5	71.1	5.3	13.2	0.405	0.314	98.3	26	
5	87.6	123.1	105.3	22.4	33.8	0.662	0.116	169.3	29	
6	26.5	173.8	118.1	30.9	2.3	13.703	0.379	132.0	33	
7	25.5	188.1	129.7	16.5	2.6	6.278	0.585	45.5	34	
8	28.3	179.6	131.6	14.4	8.3	1.740	0.237	92.2	31	

****

wood07

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	151.1	52.8	29.0	108.2	17.3	6.256	0.081	328.9	93	rsn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	34.5	150.8	0.0	19.7	5.3	3.746	0.332	72.9	37	rsn

****

wood08

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	57.7	0.0	26.7	9.6	8.6	1.109	0.696	25.7	51	pkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	31.7	156.7	7.9	18.7	3.8	4.961	0.451	51.3	29	
2	10.2	123.7	10.9	10.7	1.5	7.088	0.638	45.2	15	
3	30.5	190.3	13.9	14.4	3.8	3.827	0.563	53.3	8	
4	1507.7	89.5	18.8	80.0	38.0	2.105	0.497	89.4	63	lkn
5	45.5	81.0	22.6	8.5	9.8	0.872	0.546	31.1	54	pkn
6	37.3	171.6	29.7	33.0	2.3	14.648	0.500	102.4	26	
7	11.4	71.4	38.4	9.6	3.4	2.835	0.352	51.2	21	
8	33.5	89.0	53.0	8.5	5.3	1.620	0.746	23.0	57	pkn
9	898.4	100.2	64.3	51.7	42.1	1.228	0.413	78.2	48	lkn
10	111.6	148.7	101.1	42.6	12.0	3.544	0.218	127.0	47	pkn
11	53.1	163.6	110.9	33.6	9.4	3.572	0.168	116.9	40	

****

wood09

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	682.0	30.9	26.7	207.9	9.4	22.114	0.349	245.5	73	brk
2	373.6	164.2	41.4	74.6	12.4	6.014	0.403	168.2	35	
3	4333.4	0.0	54.9	238.8	57.5	4.151	0.315	217.0	107	brk
4	27.1	97.0	77.1	18.1	3.0	6.025	0.496	57.4	22	
5	12.0	18.7	97.8	4.8	3.4	1.418	0.741	21.6	34	

****

wood10

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	81.4	174.3	120.3	44.2	11.7	3.795	0.158	207.1	34	
2	167.5	196.7	121.4	42.1	15.8	2.666	0.252	139.9	29	
3	1164.6	54.9	135.4	183.9	33.1	5.557	0.191	184.0	114	brk
4	17.4	29.3	144.0	20.3	2.3	8.978	0.382	97.3	40	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	3051.2	69.3	51.9	95.9	56.8	1.690	0.560	68.6	54	lkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	228.9	172.7	60.2	32.0	12.0	2.658	0.595	38.6	29	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	74.2	206.8	88.7	32.0	13.2	2.430	0.176	103.8	39	

****

wood11

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1250.9	68.2	144.4	170.6	20.7	8.248	0.355	383.9	63	
2	11.0	138.6	147.8	6.4	3.8	1.701	0.458	53.0	17	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	81.2	119.4	74.4	11.7	9.0	1.299	0.767	22.8	38	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	131.5	108.7	126.3	13.9	13.2	1.053	0.721	22.7	63	pkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	211.8	155.6	20.7	19.7	17.7	1.116	0.608	28.0	63	skn

****

wood12

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	66.7	53.8	104.9	9.6	9.0	1.063	0.771	21.2	85	pkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	64.9	117.3	51.5	22.9	5.3	4.354	0.538	41.5	119	brk
2	10.8	99.1	56.8	11.7	1.9	6.237	0.491	54.0	12	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	48.5	177.5	29.3	9.6	10.5	0.911	0.480	57.5	33	
2	25.7	189.2	41.0	6.4	6.8	0.945	0.593	30.0	40	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1623.7	170.6	120.3	68.2	45.9	1.487	0.519	66.1	60	lkn
2	12.8	229.2	140.2	9.6	2.6	3.645	0.508	39.1	25	

****

wood13

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	118.8	41.0	54.1	12.3	13.2	0.932	0.737	24.3	35	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	334.9	126.3	127.1	49.0	34.6	1.418	0.197	121.2	99	lkn

****
wood14

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	208.0	58.6	136.9	20.3	16.2	1.253	0.635	25.3	51	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	296.8	71.4	103.0	29.3	15.8	1.856	0.641	32.7	43	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	11.2	165.8	3.8	10.1	1.9	5.387	0.589	41.1	26	
2	28.1	178.0	4.5	7.5	8.6	0.863	0.435	57.9	24	
3	11.6	188.7	7.9	3.7	4.1	0.902	0.753	22.3	22	
4	15.8	175.9	12.4	6.9	6.4	1.084	0.357	65.6	16	
5	223.1	154.6	18.8	29.3	29.3	1.000	0.259	111.3	55	lkn
6	17.6	145.5	35.0	6.4	5.6	1.134	0.489	38.2	40	

****
wood15

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	16.0	108.7	48.1	8.0	3.0	2.658	0.667	26.5	22	
2	10.2	117.3	51.1	5.3	3.4	1.575	0.567	28.3	26	
3	82.4	97.5	51.5	20.8	12.4	1.675	0.319	107.3	28	
4	10.8	87.9	63.5	6.4	3.0	2.126	0.563	29.6	40	
5	42.7	92.7	63.5	17.1	6.4	2.668	0.392	45.1	38	
6	10.8	105.5	67.3	9.6	3.0	3.189	0.375	50.1	45	
7	34.1	63.4	69.2	9.6	5.6	1.701	0.630	37.6	18	
8	15.2	62.9	76.0	6.9	3.4	2.048	0.650	30.3	18	
9	63.5	101.3	84.2	11.7	7.5	1.559	0.720	25.6	28	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	221.7	72.0	104.9	19.7	15.4	1.279	0.729	23.1	36	skn

****
wood16

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	92.4	89.0	30.8	11.2	19.6	0.572	0.422	114.8	41	skn

****
wood17

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	152.5	56.5	52.3	11.7	17.7	0.664	0.736	25.0	38	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	116.4	139.1	75.2	14.4	14.7	0.981	0.552	54.5	44	mkn

****
wood18

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	2614.5	75.2	0.0	138.0	71.1	1.943	0.267	257.3	113	lkn
2	30.9	119.4	28.2	13.9	4.9	2.835	0.456	39.5	62	

****
wood19

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	490.6	16.0	96.3	41.6	21.1	1.974	0.560	42.9	42	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	366.3	46.9	120.3	21.3	22.6	0.945	0.762	21.9	49	mkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	

1	51.7	85.8	0.0	19.7	16.9	1.166	0.155	106.8	53	
2	246.1	101.3	0.0	34.1	30.1	1.134	0.240	125.1	28	lkn
3	146.1	102.9	11.3	30.9	26.7	1.158	0.177	170.0	30	
4	19.4	99.7	16.5	6.4	6.0	1.063	0.505	32.3	38	

****

wood21

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	32.3	30.4	25.2	8.5	5.6	1.512	0.671	23.9	31	pkn
2	97.4	41.6	27.8	15.5	13.2	1.175	0.479	40.3	52	skn
3	15.2	37.8	45.9	4.8	4.1	1.160	0.768	21.1	65	pkn
4	20.4	53.8	60.9	8.5	6.8	1.260	0.354	48.0	24	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	61.3	71.4	89.5	14.9	18.0	0.827	0.228	75.5	60	
2	32.9	69.8	115.8	13.3	13.9	0.958	0.177	93.8	19	
3	32.5	117.8	125.6	6.9	7.1	0.970	0.656	30.2	16	
4	879.0	82.6	130.1	60.8	38.4	1.584	0.377	92.2	56	lkn
5	32.7	136.4	152.7	8.5	9.0	0.945	0.424	88.3	31	

****

wood22

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	30.3	103.4	24.1	8.0	8.6	0.924	0.438	38.3	57	
2	2684.9	0.0	30.5	238.8	28.6	8.356	0.393	116.3	76	pth
3	10.8	28.8	58.3	3.2	4.9	0.654	0.692	26.7	41	

****

wood24

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	307.2	58.6	83.5	25.6	17.7	1.448	0.680	35.7	33	skn
2	611.2	70.4	104.5	32.0	27.4	1.165	0.696	40.6	38	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	239.3	110.9	18.4	25.6	21.4	1.194	0.436	60.2	51	lkn
2	17.0	98.1	25.6	9.1	5.6	1.607	0.333	48.2	38	
3	149.7	98.6	38.4	22.9	25.6	0.896	0.255	78.4	47	

****

wood26

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	112.8	85.3	57.2	14.9	11.7	1.280	0.649	25.6	106	sho
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	105.8	109.8	24.1	14.9	10.9	1.369	0.650	39.3	24	skn

****

wood31

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	18.0	14.9	136.9	4.8	5.3	0.911	0.714	25.6	48	pkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	762.4	34.1	12.0	21.9	56.8	0.385	0.614	55.1	54	lkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	14.2	45.8	111.7	4.3	4.1	1.031	0.807	20.3	36	pkn

****

wood32

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	362.5	10.7	91.4	24.5	21.1	1.164	0.702	23.5	52	mkn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	426.9	140.2	0.0	24.5	21.1	1.164	0.827	24.0	58	mkn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	670.2	142.8	72.2	27.7	43.6	0.635	0.554	48.8	48	mkn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	101.6	173.8	51.5	12.8	10.5	1.215	0.754	21.3	48	skn

****

wood38

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	11.4	28.8	24.1	10.1	1.9	5.387	0.600	40.4	9	
2	742.3	32.5	24.1	46.9	32.3	1.451	0.489	79.3	54	skn
3	34.5	20.3	51.5	6.9	6.0	1.152	0.827	19.6	59	pkn
4	134.3	21.9	61.3	19.7	11.7	1.692	0.584	36.3	55	pkn
5	79.6	18.1	78.2	10.7	10.2	1.050	0.735	23.2	63	pkn
6	217.6	42.6	89.1	33.0	16.9	1.953	0.389	56.6	36	
7	78.0	48.0	103.8	14.4	13.2	1.094	0.412	107.0	16	
8	26.9	79.4	110.2	8.0	6.0	1.329	0.558	36.6	20	
9	1215.5	56.5	112.0	45.8	56.4	0.813	0.470	115.8	37	lkn
10	10.0	86.3	119.2	5.3	5.6	0.945	0.333	58.3	14	
11	11.0	49.6	137.2	7.5	6.4	1.167	0.231	69.9	26	

FEATURE DESCRIPTION - TRAINING SET

ADAPTIVE THRESHOLD METHOD

This list contains all the features listed in appendix 1A and describes how well they are extracted using the adaptive threshold method to define the feature area extent. The features are listed in the catagories of table 9.2 and list DY, the feature size as determined by a Quality Control Officer, and DY', the size of the blob as determined by the discrimination process (both im millimetres).

TRAINING SET

LARGE KNOTS

IMAGE	DY	DY'	COMMENT
WOOD01	37.2	32.0	bark only, 1 extra small blob
WOOD03	33.5	30.8	2 extra small blobs, well defined
WOOD04	33.1	35.7	well defined
WOOD06	144.4	144.4	7 extra blobs, spike knot, well defined
WOOD08	30.1	38.0	3 pin knots and 6 g-ring blobs, well defined
WOOD08	42.1	42.1	same feature area as above, well defined
WOOD10	44.0	56.8	well defined but extended by g-ring slightly
WOOD12	55.3		2 bark blobs, feature undefined
WOOD12	30.8	45.9	1 extra blob, extended by g-ring
WOOD13	32.7	34.6	bark only, defines feature size
WOOD14	47.8	29.3	6 bark blobs, feature poorly defined
WOOD15	37.2		8 bark blobs, feature undefined
WOOD18	78.2	71.1	bark only, 1 extra blob, extent defined well
WOOD19	41.0	30.1	4 bark blobs, feature poorly defined
WOOD21	107.2	38.4	5 blobs, feature poorly defined
WOOD24	47.4	21.4	3 bark blobs, feature poorly defined
WOOD31	52.3	56.8	bark and centre, passable definition
WOOD38	77.8	56.4	6 blobs, passably defined

MEDIUM KNOTS

IMAGE	DY	DY'	COMMENT
WOOD02	28.6	28.6	hole, well defined
WOOD13	25.6		not detected, classed clear
WOOD15	25.2		not detected, classed clear
WOOD16	24.8		classed clear
WOOD17	27.1	14.7	bark only, does not define feature size
WOOD19	22.9	22.6	half defined, defines feature size
WOOD32	20.7	21.1	defined well
WOOD32	21.4	21.1	defined well
WOOD32	27.1	43.6	defined well but g-ring extends feature size

SMALL KNOTS

IMAGE	DY	DY'	COMMENT - defined well unless otherwise
WOOD01	11.3	11.7	hole
WOOD11	13.5	9.0	

WOOD11	16.5	17.7	
WOOD13	13.9	13.2	
WOOD14	14.7	16.2	
WOOD14	15.4	15.8	
WOOD15	12.4	7.5	
WOOD15	14.3	15.4	
WOOD16	17.3	19.6	bark only, extent defined well
WOOD17	18.0	17.7	half defined, size defined well
WOOD17	13.2		classed clear
WOOD19	19.6	21.1	
WOOD21	12.4	13.2	
WOOD24	13.5	17.7	
WOOD24	19.9	27.4	
WOOD26	11.3	11.7	hole
WOOD26	16.5	10.9	
WOOD32	10.5	10.5	
WOOD38	16.5	32.3	

PIN KNOTS

IMAGE	DY	DY'	COMMENT - well defined unless otherwise
WOOD02	6.0	7.9	
WOOD08	7.9	8.6	
WOOD08	7.5	9.8	
WOOD08	5.6	5.3	
WOOD08	6.4	12.0	size extended by g-ring
WOOD11	9.4	13.2	
WOOD12	8.6	9.0	hole
WOOD14	7.1		classed clear
WOOD21	5.3	5.6	
WOOD21	4.1	4.1	
WOOD31	5.6	5.3	
WOOD31	4.9	4.1	
WOOD38	6.0	6.0	
WOOD38	8.3	11.7	
WOOD38	9.4	10.2	

PITH

IMAGE	COMMENT
WOOD05	well defined
WOOD22	well defined

RESIN AND BARK

IMAGE	COMMENT
WOOD07	resin, well defined
WOOD07	resin, well defined
WOOD09	bark, well defined
WOOD09	bark, well defined but extended by false positive area
WOOD10	bark, well defined, 2 extra g-ring blobs
WOOD11	bark, extended by g-rings
WOOD12	bark, small piece defined well

APPENDIX 4B - FEATURE DISCRIMINATION

TEST SET - ADAPTIVE THRESHOLD METHOD

The photographs presented in this appendix are the processed results of the test set of the WOOD library of images. For the purposes of display the original images are reduced to one-quarter size by displaying only every second pixel in every second row of each image. In this way four images are displayed on one screen and photographed.

Each figure is the result of processing using the adaptive threshold method to determine the feature area extent. Each feature area is processed using the modification of the adaptive threshold method for feature areas. The grid lines are the boundaries of the local areas of size 64 * 64 pixels. The white rectangles are the feature areas and each coloured blob is separate from every other. Vacant areas that can be inferred by the absence of grid lines are feature areas that have been detected and classed as clear.

Following the figures is a list of the location and size of every blob with various descriptive shape measures. This is followed by a description of how well each feature is described.

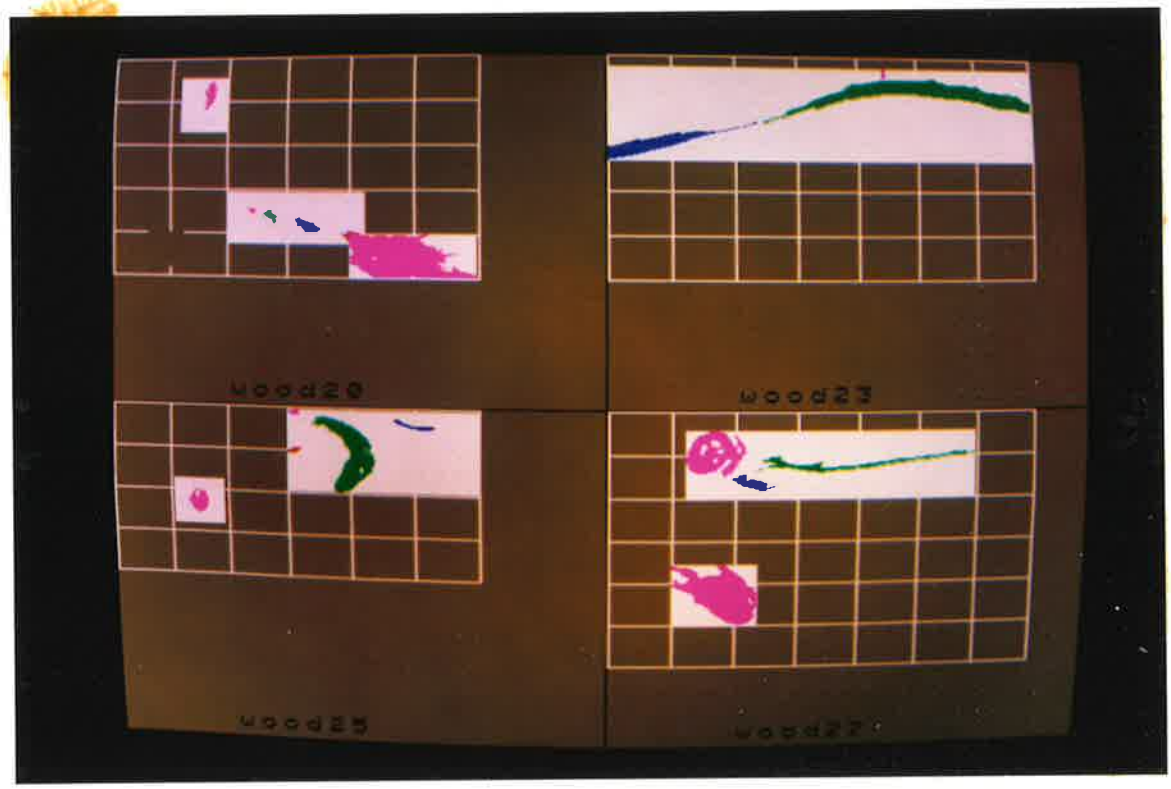

Figure A4B.1 Test set - WOOD20, WOOD23, WOOD25, WOOD27.

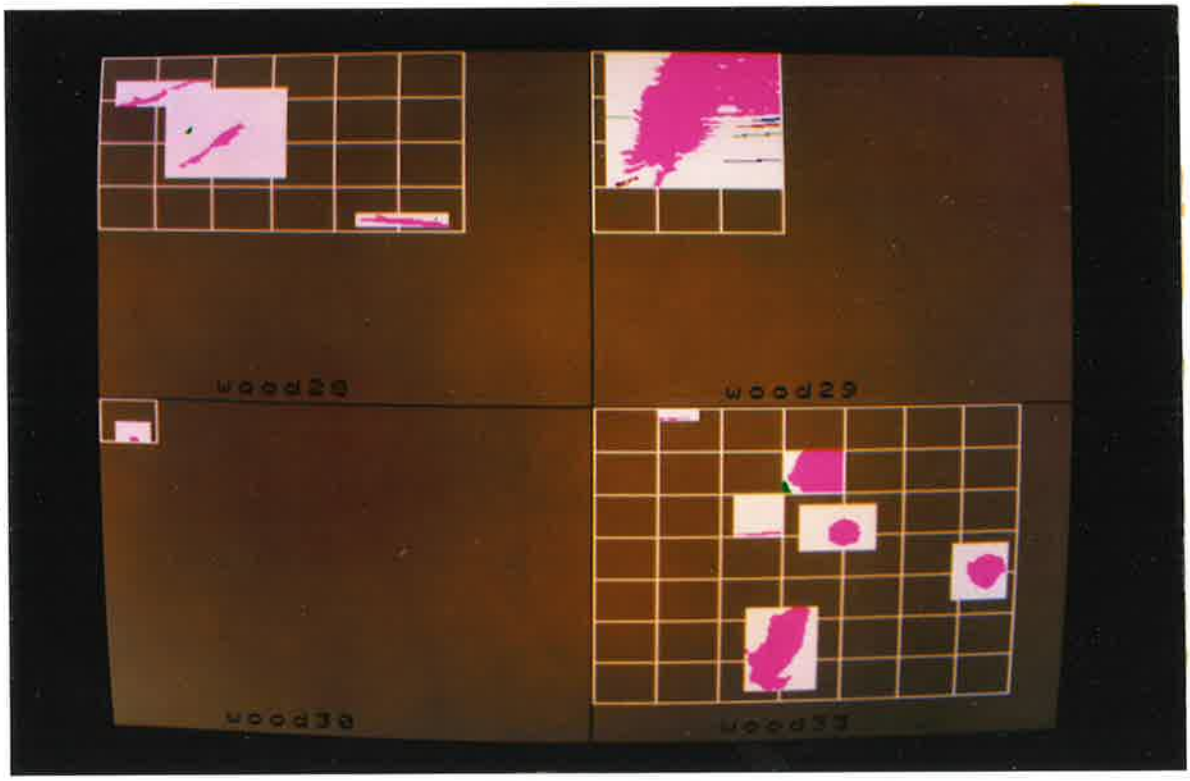

Figure A4B.2 Test set - WOOD28, WOOD29, WOOD30, WOOD33.

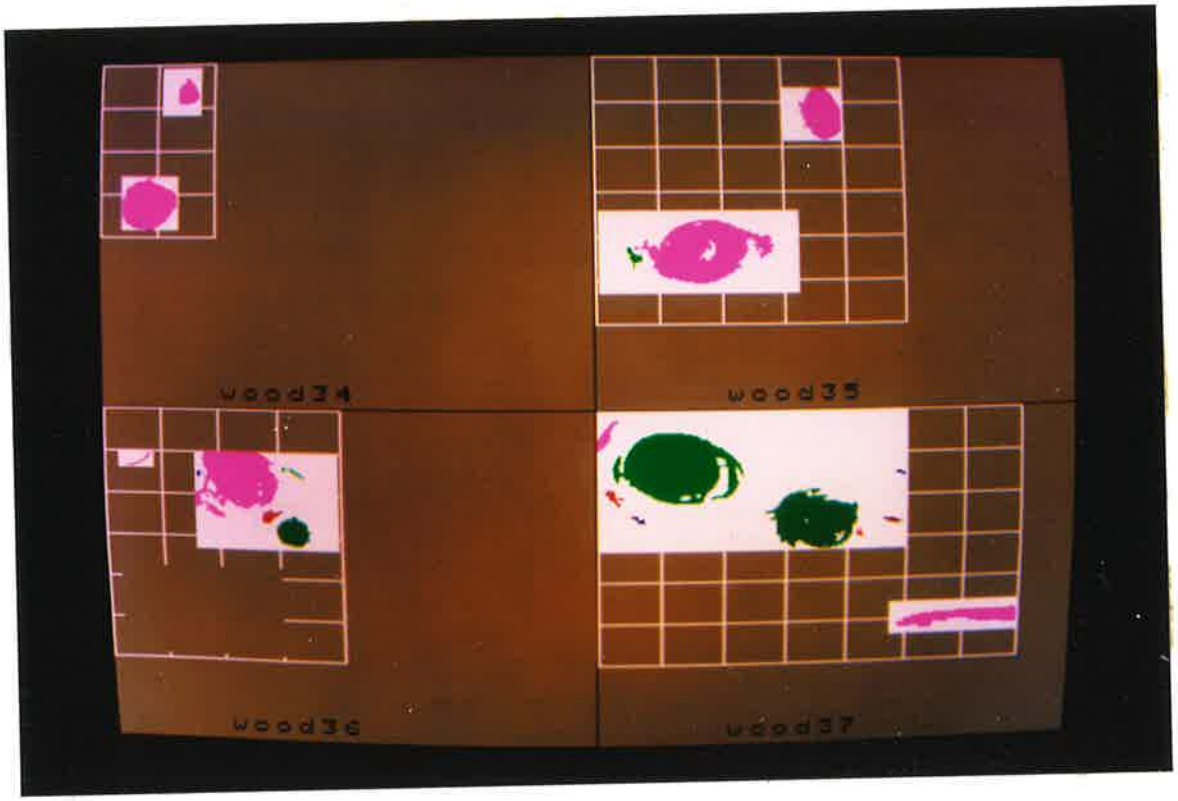

Figure A4B.3 Test set - WOOD34, WOOD35, WOOD36, WOOD37.

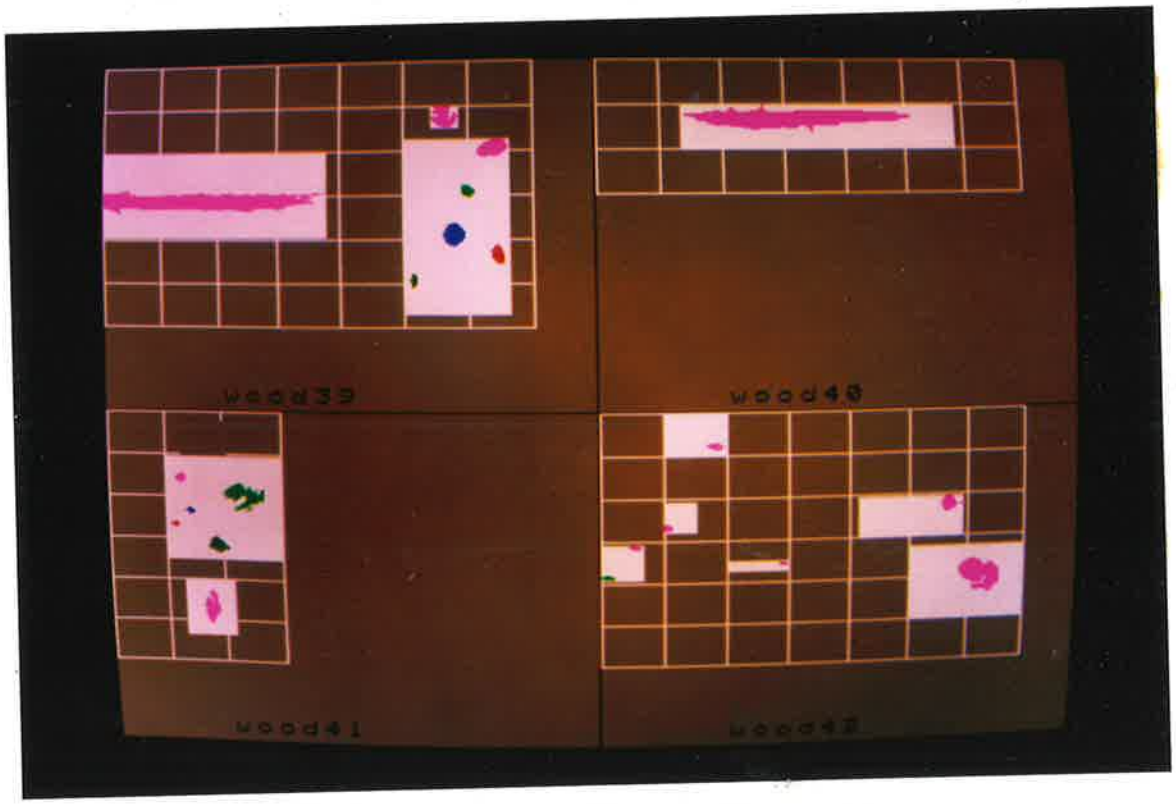

Figure A4B.4 Test set - WOOD39, WOOD40, WOOD41, WOOD42.

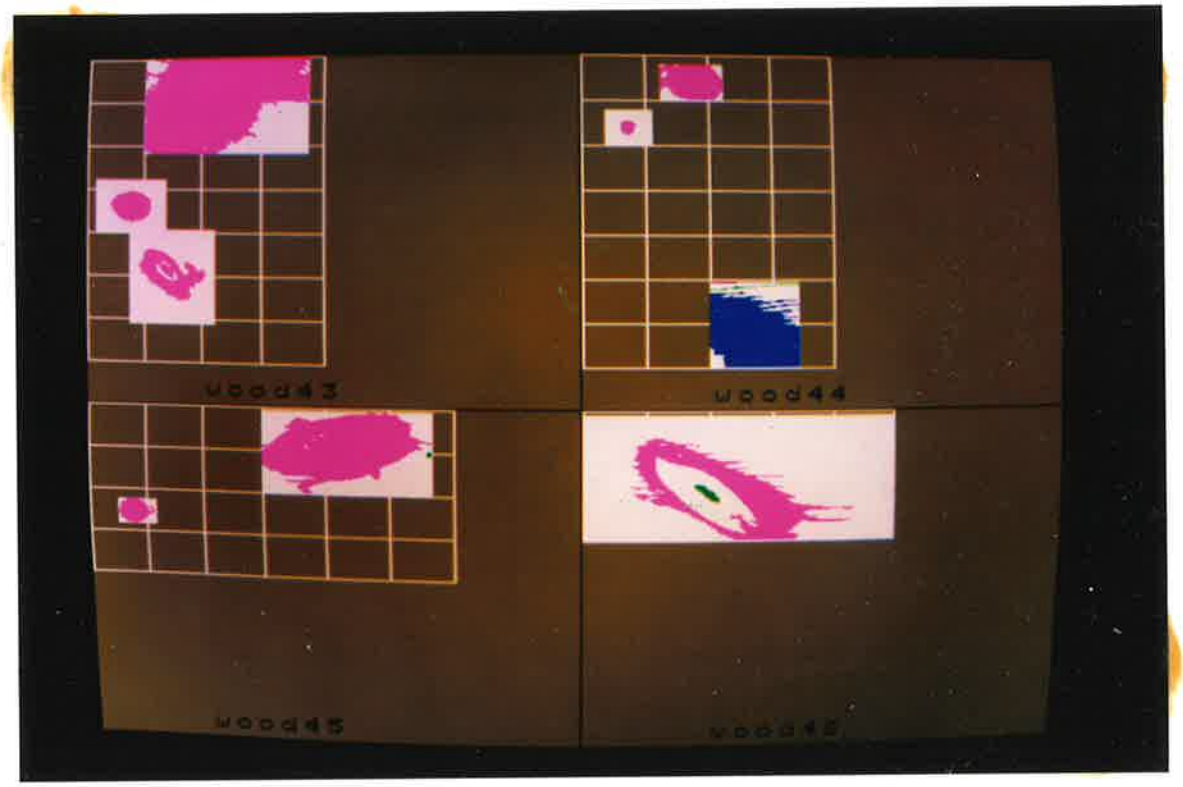

Figure A4B.5 Test set - WOOD43, WOOD44, WOOD45, WOOD46.

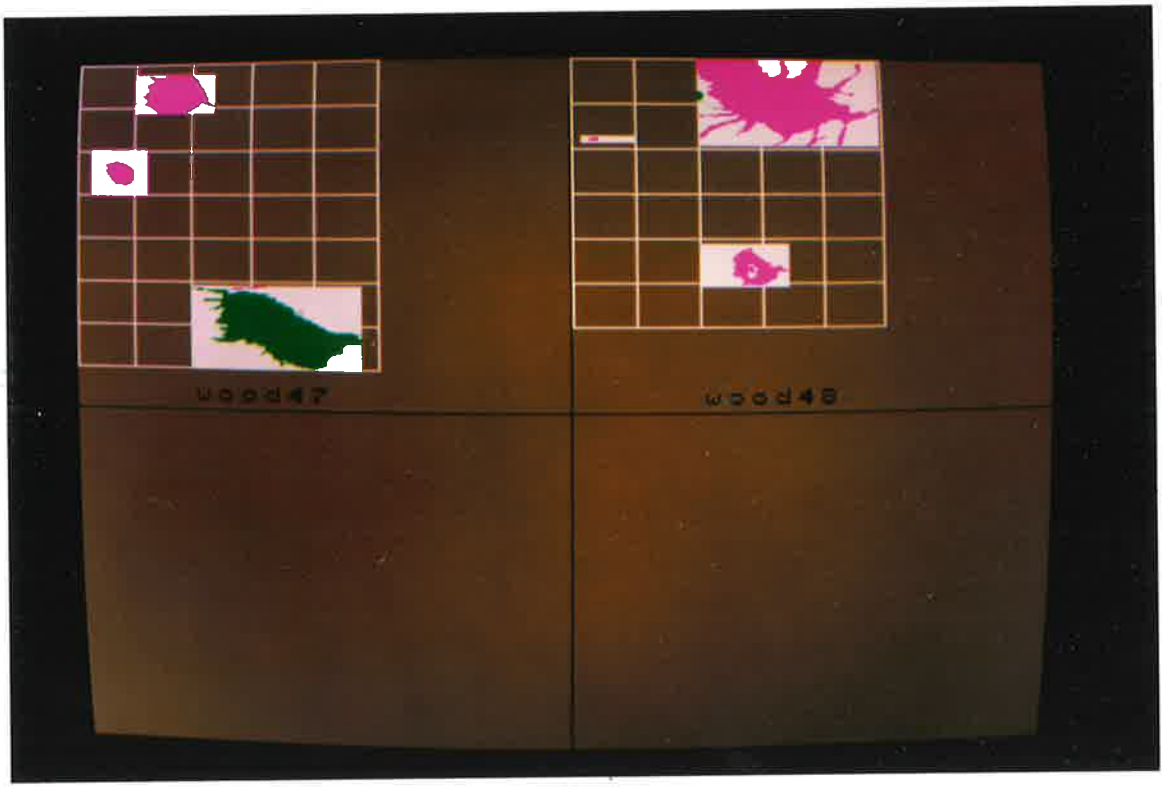

Figure A4B.6 Test set - WOOD47, WOOD48.

BLOB DESCRIPTION - TEST SET

ADAPTIVE THRESHOLD METHOD

Presented below is a list of all the blobs identified in the previous figures. The measures are:

area in mm^2, of the blob.

x,y in mm, of the top left corner of a box that encloses the blob. The origin is the top left corner of each image.

dx,dy in mm, the size of the blob. The significant dimension for the grading rules is dy which is the size of the feature.

len/wid the ratio of the length to the width (dx/dy) is a measure of the shape of the blob.

propn the proportion of the enclosing box of the blob that contains blob pixels.

p^2/a ratio of the square of the perimeter of the blob to the area. This is a measure of the shape.

coldiff the difference in average grey levels between the blob pixels and the background pixels in the box that encloses the blob.

The blobs that define features are marked with the following code:

Large knot(hole)	lkn(lho)
Medium knot(hole)	mkn(mho)
Small knot(hole)	skn(sho)
Pin knot(hole)	pkn(pho)
Needle trace	ntr
Pith	pth
Resin	rsn
Bark	brk
Growth ring	grn

****

wood20

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	68.3	53.8	12.8	7.5	15.8	0.473	0.580	42.2	51	lkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	11.0	81.0	82.3	4.3	3.4	1.260	0.764	21.0	28	
2	26.3	88.5	83.8	9.1	6.8	1.339	0.428	44.1	28	
3	49.7	106.6	88.0	13.9	7.5	1.843	0.477	37.2	20	
4	66.7	132.2	94.8	12.8	7.5	1.701	0.694	26.5	38	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1107.3	136.4	96.3	67.2	24.1	2.791	0.685	43.8	80	mkn

****

wood23

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	10.6	150.8	5.6	3.2	5.3	0.608	0.631	30.2	34	
2	1011.3	82.6	10.9	156.2	25.2	6.199	0.257	130.0	76	pth

3	348.9	0.0	35.0	81.5	21.1	3.873	0.203	105.2	79	pth

****

wood25

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	111.6	44.8	47.8	11.7	13.5	0.866	0.703	24.2	82	sho
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	16.4	102.3	1.5	7.5	3.0	2.481	0.732	28.1	8	
2	567.4	115.1	5.3	34.6	42.1	0.823	0.389	59.4	77	mkn
3	41.3	159.4	5.6	22.4	6.0	3.721	0.307	65.3	66	
4	16.2	102.3	22.2	8.5	3.4	2.520	0.563	30.9	48	

****

wood27

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	533.9	41.6	10.2	33.0	25.6	1.292	0.632	53.1	25	mkn
2	297.4	78.9	22.6	124.7	10.9	11.438	0.219	286.5	51	rsn
3	86.4	66.6	34.6	23.5	9.0	2.599	0.408	53.6	29	rsn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	929.7	34.1	85.7	46.9	34.2	1.371	0.579	47.2	53	mkn

****

wood28

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	142.7	9.1	13.2	58.1	14.7	3.962	0.167	135.0	57	rsn
2	62.5	50.1	22.6	17.1	5.3	3.240	0.696	27.1	36	rsn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	164.3	49.0	36.5	37.8	25.2	1.502	0.172	98.4	56	lkn
2	11.2	52.2	39.5	4.8	3.8	1.276	0.622	25.8	26	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	129.9	149.2	87.2	47.4	6.8	7.009	0.404	74.7	81	rsn

****

wood29

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	3346.6	14.9	0.0	87.4	72.2	1.211	0.530	88.3	74	lho
2	14.2	7.5	33.8	13.9	1.5	9.214	0.683	50.7	20	
3	21.6	86.3	34.6	16.0	2.3	7.088	0.600	48.0	14	
4	27.9	77.3	37.2	25.1	2.6	9.518	0.422	83.9	23	
5	22.0	75.7	42.5	26.7	2.3	11.813	0.367	114.0	22	
6	28.3	70.4	56.8	32.0	1.9	17.011	0.470	123.6	22	
7	22.2	11.2	66.6	13.3	5.6	2.363	0.296	86.5	9	

****

wood30

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	17.2	17.6	20.7	5.9	3.4	1.733	0.869	18.6	36	skn

****

wood33

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	18.0	34.1	4.9	18.1	1.5	12.049	0.662	67.6	25	rsn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	38.9	75.7	69.2	26.7	3.0	8.860	0.485	69.4	42	rsn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	998.0	83.1	113.2	36.2	48.9	0.741	0.563	47.0	51	lkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	657.3	105.0	24.1	31.4	24.1	1.307	0.869	18.8	64	mkn
2	25.1	102.3	40.6	5.3	7.5	0.709	0.625	28.8	21	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	226.5	128.5	63.2	18.7	15.8	1.181	0.769	22.7	35	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	339.5	210.0	85.7	24.0	20.3	1.181	0.697	23.1	56	skn

****

wood34

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	716.5	12.8	62.0	33.0	30.1	1.099	0.721	24.2	72	mkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	129.7	46.4	7.9	13.3	14.3	0.933	0.681	24.5	62	skn

****

wood35

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1490.2	20.8	87.6	76.2	34.6	2.203	0.565	53.7	35	lkn
2	49.1	16.5	101.9	9.1	12.4	0.730	0.437	45.9	26	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	491.0	112.5	16.5	24.0	28.2	0.851	0.726	28.0	49	mkn

****

wood36

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	23.4	10.1	24.1	18.7	9.0	2.067	0.139	119.0	69	rsn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	1193.0	54.4	24.1	48.5	37.2	1.303	0.661	66.2	31	mkn
2	24.9	103.9	32.3	13.9	6.4	2.168	0.281	65.3	24	
3	11.0	54.9	35.0	5.3	4.1	1.289	0.500	49.2	7	
4	45.9	92.7	56.4	12.3	7.5	1.630	0.498	40.2	33	
5	234.5	100.7	60.2	19.2	17.3	1.109	0.707	25.3	46	skn

****

wood37

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	73.1	0.0	4.9	11.7	17.3	0.678	0.361	53.7	25	
2	1923.3	7.5	11.7	71.4	40.2	1.775	0.669	61.8	43	lkn
3	11.8	164.2	34.6	6.4	3.4	1.890	0.546	29.9	28	
4	39.1	3.2	43.6	11.7	9.8	1.199	0.341	64.3	23	
5	1180.8	90.1	44.0	52.8	34.2	1.542	0.654	45.5	51	mkn
6	19.2	17.6	58.3	8.5	4.9	1.745	0.462	45.4	25	
7	21.0	157.8	61.7	11.7	2.6	4.455	0.682	32.0	31	
8	12.0	141.2	66.6	4.8	4.9	0.981	0.513	35.3	26	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	473.8	165.8	115.8	73.0	10.5	6.936	0.616	57.9	51	rsn

****

wood39

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	997.4	0.0	68.1	129.5	13.2	9.842	0.585	80.8	71	pth

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	113.0	185.5	24.1	14.4	11.7	1.235	0.674	37.8	44	skn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	121.4	209.5	41.7	17.1	10.2	1.680	0.701	23.8	46	skn
2	40.7	200.4	66.6	8.0	7.1	1.119	0.712	22.8	51	pkn
3	112.2	191.9	87.2	11.7	13.2	0.891	0.727	23.2	61	skn
4	49.5	216.4	99.3	6.9	10.2	0.683	0.704	27.2	33	skn
5	22.8	173.8	114.7	4.3	7.9	0.540	0.679	31.6	61	pkn

****

wood40

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	906.4	47.4	24.1	128.5	15.4	8.332	0.458	81.2	69	pth

****

wood41

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	28.3	40.0	34.6	6.9	6.0	1.152	0.678	23.9	50	pkn
2	205.2	68.8	40.6	24.5	16.2	1.516	0.518	72.3	37	skn
3	15.2	45.3	54.9	6.9	3.4	2.048	0.650	27.8	27	pkn
4	11.6	36.2	62.8	6.4	3.0	2.126	0.604	27.6	36	pkn
5	74.8	59.7	71.4	12.8	9.4	1.361	0.622	26.8	24	skn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	120.6	55.4	102.6	10.1	19.9	0.508	0.598	58.7	39	mkn

****

wood42

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	17.0	14.9	74.4	7.5	3.0	2.481	0.759	22.8	34	ntr
2	18.8	0.0	91.0	8.0	3.0	2.658	0.783	24.5	29	ntr

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	30.1	57.0	16.9	11.2	4.5	2.481	0.595	30.8	58	ntr

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	17.8	34.1	62.4	5.3	4.5	1.181	0.742	21.8	46	ntr

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	11.2	96.5	84.6	5.9	2.6	2.228	0.727	23.1	45	ntr

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	66.7	189.2	49.6	13.3	8.6	1.541	0.579	39.0	52	pkn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	321.1	200.4	86.5	25.1	19.6	1.281	0.655	50.3	54	skn

****

wood43

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	319.5	13.3	72.9	25.1	16.9	1.481	0.754	21.2	48	skn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	3876.3	34.1	0.0	94.9	52.3	1.815	0.782	38.8	56	lkn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	592.0	30.9	105.3	40.0	28.2	1.418	0.525	51.0	31	mkn

****

wood44

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	52.5	20.8	35.0	8.5	7.9	1.080	0.780	20.9	50	pkn

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	551.7	42.1	4.9	34.1	19.2	1.779	0.843	22.0	60	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	13.8	68.2	121.4	16.5	1.1	14.648	0.742	67.0	25	
2	41.1	68.2	122.6	48.5	3.8	12.900	0.225	203.0	43	
3	1691.2	68.2	127.1	48.5	41.4	1.173	0.843	73.6	83	lkn

****

wood45

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	214.4	17.1	53.8	21.9	14.3	1.529	0.687	31.6	55	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	2490.3	102.3	1.5	91.1	46.6	1.955	0.586	45.0	66	lkn
2	10.0	189.2	22.9	4.3	3.4	1.260	0.694	25.9	44	

****

wood46

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	2201.5	25.1	14.7	123.1	57.5	2.140	0.311	151.0	91	lkn
2	65.5	59.7	40.2	14.4	10.5	1.367	0.433	38.4	45	

****

wood47

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	170.1	16.5	53.4	17.6	13.2	1.337	0.735	22.4	50	skn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	671.0	34.1	5.6	48.0	22.2	2.162	0.631	39.6	66	mkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	22.0	90.6	122.6	19.7	2.3	8.742	0.495	70.4	42	
2	1872.4	68.2	124.1	93.3	44.4	2.102	0.452	89.8	109	lho

****

wood48

no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	12.4	8.5	41.7	6.9	2.6	2.633	0.681	25.8	63	pkn
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	2522.1	68.2	0.0	100.7	46.6	2.161	0.537	201.9	69	mho
2	15.2	68.2	16.9	4.8	5.3	0.911	0.603	30.3	17	
no	area	x	y	dx	dy	len/wid	propn	p^2/a	coldiff	
1	337.5	85.8	101.5	31.4	20.3	1.549	0.529	40.1	70	skn

FEATURE DESCRIPTION - TEST SET

ADAPTIVE THRESHOLD METHOD

This list contains all the features listed in appendix 1A and describes how well they are extracted using the adaptive threshold method to define the feature area extent. The features are listed in the catagories of table 9.2 and list DY, the feature size as determined by a Quality Control Officer, and DY', the size of the blob as determined by the discrimination process (both im millimetres).

TEST SET

LARGE KNOTS

IMAGE	DY	DY'	COMMENT
WOOD20	34.2	15.8	bark only, does not define feature size
WOOD28	40.2	25.2	bark only, does not define feature size well
WOOD29	51.1	72.2	hole, 8 extra g-ring blobs, size defined
WOOD33	42.1	48.9	defined well
WOOD35	35.0	34.6	1 extra resin blob, defined well
WOOD37	40.2	40.2	some resin noise, defined well
WOOD43	55.3	52.3	size defined
WOOD44	35.7	41.4	2 g-ring blobs, size defined
WOOD45	32.7	46.6	1 extra blob, extended by g-ring
WOOD46	51.1	57.5	bark, plus pith blob, size defined well
WOOD47	39.9	44.4	hole, defined well

MEDIUM KNOTS

IMAGE	DY	DY'	COMMENT
WOOD20	21.4		classed clear
WOOD20	22.9	24.1	size defined well
WOOD25	26.3	42.1	bark only, defines large size, 3 extra blobs
WOOD27	24.1	25.6	2 extra resin blobs, well defined
WOOD27	29.0	34.2	well defined, size extended by g-ring
WOOD33	24.1	24.1	1 noise blob, size defined well
WOOD34	27.8	30.1	defined well
WOOD35	28.6	28.2	defined well
WOOD36	26.7	37.2	defined well, size extended by g-ring
WOOD36	29.7		classed clear
WOOD37	26.3	34.2	good, extended by g-ring, 3 extra blobs
WOOD41	21.1	19.9	bark only, size defined well
WOOD43	28.2	28.2	defined well
WOOD47	20.7	22.2	defined well
WOOD48	20.3	46.6	hole, extended by g-ring

SMALL KNOTS

IMAGE	DY	DY'	COMMENT - defined well unless otherwise
WOOD20	17.3		classed as clear
WOOD25	13.2	13.5	hole
WOOD30	12.0	3.4	defined under-size

WOOD33	18.0	15.8	
WOOD33	19.2	20.3	
WOOD34	14.3	14.3	
WOOD36	14.7	17.3	
WOOD39	15.0	11.7	
WOOD39	10.2	10.2	
WOOD39	12.8	13.2	
WOOD39	11.3	10.2	
WOOD41	10.2		classed as clear
WOOD41	12.8	16.2	
WOOD41	13.2	9.4	
WOOD42	17.3	19.6	
WOOD43	17.3	16.9	
WOOD44	16.2	19.2	
WOOD45	14.3	14.3	
WOOD47	12.8	13.2	
WOOD48	15.8	20.3	

PIN KNOTS

IMAGE	DY	DY'	COMMENT - defined well unless otherwise
WOOD39	7.5	7.1	
WOOD39	8.3	7.9	
WOOD41	6.8	6.0	
WOOD41	4.5	3.4	
WOOD41	3.8	3.0	
WOOD41	4.1		classed as clear
WOOD42	5.6	3.0	
WOOD44	8.3	7.9	
WOOD48	3.0	2.6	

NEEDLE TRACE

IMAGE	DY	DY'	COMMENT
WOOD42	2.6	3.0	needle trace
WOOD42	4.1	4.5	needle trace
WOOD42	2.3	4.5	needle trace
WOOD42	2.3	2.6	needle trace
WOOD42	3.8	8.6	needle trace

PITH

IMAGE	COMMENT
WOOD23	defined well
WOOD23	defined well
WOOD39	defined well
WOOD40	defined well

RESIN AND BARK

IMAGE	COMMENT
WOOD27	resin, defined well
WOOD28	resin, defined well
WOOD28	resin, defined well
WOOD33	resin, defined well
WOOD33	resin, defined well
WOOD33	resin, classed as clear
WOOD37	resin, threshold low, defined over-size

APPENDIX 5 - GRADING RULES FOR SAWN BOARDS

The following text is reprinted from the **Australian Standard Specification, Sawn Boards for Radiata Pine (metric units), AS 1489 - 1973**, published by the Standards Association of Australia.

SECTION 1. SCOPE AND GENERAL PROVISIONS

1.1 SCOPE. This specification applies to radiata pine boards, graded on face and/or edge appearance and intended for end-uses where appearance is the prime consideration.

1.2 GRADES. Provision is made for the following grades of boards:

Clear grade
Joinery grade
Select grade
Standard grade

1.8 MEASUREMENT OF IMPERFECTIONS.

1.8.1 Knots and Holes. Knots and holes shall be measured as the width on the face measured between lines parallel with the edges of the board.

1.8.2 Bow and Spring. Bow and spring shall be measured by stretching a string from the edge at one end of the board to the same edge at the other end and measuring the maximum distance from the string to the piece.

1.8.2 Wane. Wane shall be measured as the amount by which the width of the face or edge of a board is deficient.

1.12 GRADE LIMITS AND INSPECTION. The grade descriptions in section 2 describe timber on the lower limit of the grade. Each parcel of boards supplied shall include a reasonable distribution of material of quality ranging above the lower limit of the grade.

As only the general features of a grade can be described, variations of up to 5 per cent of a parcel between gradings of individual inspectors shall be accepted.

1.13 GENERAL QUALITY. The imperfections as limited in a grade description shall be permitted in any board and be unlimited in occurrence except when the description limits the occurence of a particular imperfection in relation to length or to a percentage of the number of boards in a parcel.

Imperfections that can be removed by normal machining shall be permitted.

SECTION 2. GRADE DESCRIPTIONS

2.1 CLEAR GRADE. Each clear grade board shall be of sound wood, truly sawn, and shall be free from defects on both faces and both edges. The following shall be permitted:

(i) **Bow** - evenly distributed and -

(a) for thickness of 24 mm and less, not exceeding the equivalent of 46 mm in 3 m;

(b) for thicknesses over 24 mm and up to and including 40 mm, not exceeding the equivalent of 19 mm in 3 m;

(c) for thicknesses over 40 mm, not exceeding the equivalent of 10 mm in 3 m.

(ii) **Spring** - evenly distributed and -

(a) for widths of 96 mm or less, not exceeding the equivalent of 29 mm in 3 m;

(b) for widths over 96 mm and up to and including 146 mm, not exceeding the equivalent of 19 mm in 3 m;

(c) for widths over 146 mm, not exceeding the equivalent of 11 mm in 3 m.

(iii) **Sloping grain** - not exceeding 1 in 6.

2.2 JOINERY GRADE

2.2.1 General. Each joinery grade board shall be graded on one face and both edges, and shall be sound wood, truly sawn. The imperfections listed in Clauses 2.2.2 and 2.2.3 shall be permitted.

2.2.2 On the Best Face and Both Edges.

(i) **Sound knots, and partially encased knots** - not exceeding one-half of the width of the face or edge, 25mm individually, or 50 mm in aggregate whichever is the least at the worst cross-section of the board; and not more than three knots of maximum permissible size in 4 m of board. Any encasement of knots shall not exceed one-half of the perimeter of the knot, and any void associated with the encasement shall not exceed 3 m wide nor extend through the board.

(ii) **Tight encased knots** - not exceeding 10 mm; not more than one such knot in any 2 m of board; and not more than 20 per cent of the number of boards in a parcel so affected.

(iii) **Cone holes and knot holes** - not exceeding 10 mm; not more than one such hole in any 2 m length of board; not occurring

on the edges; and not more than 10 percent of the number of boards in a parcel so affected.

(iv) **Resin pockets and bark pockets** - not exceeding 6 mm wide and 50 mm long, including any void associated with the pockets but provided the void does not extend through the board.

(v) **Checks.**

(a) Surface checks - not exceeding 1 mm wide.

(b) Seasoning checks in otherwise sound knots - not exceeding 2 mm wide.

(vi) **Needle trace** - sound and slight.

(vii) **Bow** - evenly distributed and -

(a) for thicknesses of 24 mm or less, not exceeding the equivalent of 46 mm in 3 m;

(b) for thicknesses over 24 mm and up to and including 40 mm, not exceeding the equivalent of 28 mm in 3 m;

(c) for thicknesses over 40 mm, not exceeding the equivalent of 10 mm in 3 m.

(viii) **Spring** - evenly distributed and -

(a) for widths of 96 mm or less, not exceeding the equivalent of 29 mm in 3 m;

(b) for widths over 96 mm and up to and including 146 mm, not exceeding the equivalent of 19 mm in 3 m;

(c) for widths over 146 mm, not exceeding the equivalent of 11 mm in 3 m.

(ix) **Blue stain** - slight, and not occuring in more than 5 per cent of the number of boards in a parcel.

(x) **Sloping grain.**

2.2.3 **On the Back.** Wane shall not be permitted, but other imperfections, including those referred to in Clause 2.2.2 above shall be permitted, provided they do not unduly impair the strength of the board.

2.3 SELECT GRADE.

2.3.1 General. Each select grade board shall be graded on one face only and shall be of sound wood, truly sawn. The imperfections listed in Clauses 2.3.2 and 2.3.3 shall be permitted.

2.3.2 On the Face.

(i) **Sound knots** - individual width or aggregate width at worst cross-section of the board not exceeding one-half of the width of the piece, and provided that partial encasement of any knot does not exceed 50 percent of the perimeter of the know and that any void associated with the encasement does not exceed 6 mm wide not extend through the board.

(ii) **Tight encased knots** - not exceeding 20 mm; not more than two such knots in any 4 m length of board; and not more than 15 per cent of the number of boards in a parcel so affected.

(iii) **Cone holes and knot holes** - not exceeding 12 mm; not more than two such holes in any 4 m length of board and not more than 15 per cent of the number of boards in the parcel so affected.

(iv) **Resin pockets and bark pockets** - not exceeding 6 mm wide and 75 mm long, including any void associated with the pockets but provided the void does not extend through the board.

(v) **Checks.**

(a) Surface checks - not exceeding 1 mm wide.
(b) Seasoning checks - in otherwise sound knots and not exceeding 2 mm wide.

(vi) **Pith** - sound, not exceeding 10 mm wide and 300 mm long individually or in aggregate, and not more than 5 per cent of the number of boards in the parcel so affected.

(vii) **Needle trace** - sound.

(viii) **Bow** - evenly distributed, and not exceeding the equivalent of 50 mm in 3 m.

(ix) **Spring** - evenly distributed and -

(a) for widths of 146 mm or less, not exceeding the equivalent of 28 mm in 3 m;

(b) for widths over 146 mm, not exceeding the equivalent of 18 mm in 3 m.

(x) **Blue stain** - slight.

(xi) **Sloping grain and cross grain.**

2.3.3 On the Back and Edges.

(i) **Wane** - not exceeding 12 mm wide on the back and 6 mm wide on edges.

(ii) **Other imperfections** - including those described in Clause 2.3.2 above, provided they do not unduly impair the strength of the board.

2.4 STANDARD GRADE.

2.4.1 General. Each standard grade board shall be graded on one face only and shall be of sound wood truly sawn. The imperfections listed in Clauses 2.4.2 and 2.4.3 shall be permitted.

2.4.2 On the Face.

(i) **Sound knots** - individual width or aggregate width at the worst cross-section of the board not exceeding one-half of the width of the board, and provided that partial encasement of any knot does not exceed 50 per cent of the perimeter of the knot and that ny void associated with the encasement does not exceed 6 mm wide nor extend through the board.

(ii) **Tight encased knots.**

(a) Up to 10 mm wide - not more than 20 percent of the number of boards in the parcel so affected.

(b) Over 10 mm wide, but not exceeding 20 mm wide - not more than one such knot in any 4 m length of board, and not more than 20 per cent of the number of boards in the parcel so affected.

(iii) **Cone holes and knot holes** - not exceeding 12 mm; not more than two such holes in any 4 m length of board; and not more than 20 per cent of the number of boards in the parcel so affected.

(iv) **Resin pockets and bark pockets** - not exceeding 10 mm wide, including any void associated with the pockets, but provided the void does not extend through the board.

(v) **Checks.**

(a) Surface checks - not exceeding 1 mm wide.

(b) Seasoning checks - in otherwise sound knots and not exceeding 2 mm wide.

(vi) **Pith.**

(a) All lengths of board - sound, not exceeding 10 mm wide and 450 mm long individually, and not more than 20 per cent of the number of boards in the parcel so affected.

(b) Boards up to 3 m long - in addition to (a) above, aggregate length of all occurrences not exceeding 500 mm.

(c) Boards over 3 m long - in addition to (a) above, aggregate length of all occurrences not exceeding 900 mm.

(vii) **Needle trace** - sound.

(viii) **Bow** - evenly distributed and not exceeding the equivalent of 50 mm in 3 m.

(ix) **Spring** - evenly distributed and not exceeding the equivalent of 34 mm in 3 m.

(x) **Blue stain** - slight.

(xi) **Sloping grain and cross grain.**

(xii) **Plugged holes.** Holes and unsound knots plugged in accordance Appendix B shall be regarded as sound knots.

2.4.3 **On the Back and Edges.**

(a) **Wane** - not exceeding 6 mm wide on edges and 12 mm wide on the back.

(b) **Other imperfections** - including those described in Clause 2.4.2 above, proveded they do not unduly impair the strength of the board.

BIBLIOGRAPHY

BAIRD, M.L., 1983, "GAGESIGHT: a computer vision system for automatic inspection of instrument gauges", IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. PAMI-5, No. 6, Nov, pp. 618-621.

BARTLETT, S.L., BESL, P.J., COLE, C.L., JAIN, R., MUKHERJEE, D., SKIFSTAD, K.D., 1988, "Automatic solder joint inspection", IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. PAMI-10, No. 1, Jan, pp. 31-43.

BATCHELOR, B.G., 1986, "Systems considerations for automated visual inspection", Proceedings of the first IFAC workshop on digital image processing in industrial applications, Espoo, Finland, Jun., pp. 15-32.

BESL, P.J., DELP, E.J., JAIN, R., 1985, "Automatic visual solder joint inspection", IEEE Journal of Robotics and Automation, Vol. RA-1, No. 1, Mar., pp. 42-56.

BLANZ, W.E., SANZ, J.L.C., HINKLE, E.B., 1988, "Image analysis methods for solder-ball inspection in integrated circuit manufacturing", IEEE Journal of Robotics and Automation, Vol. RA-4, No. 2, Apr., pp. 129-139.

BOURKAROUBA, S., REBORDAO, J.M., WENDEL, P.L., 1985, "An amplitude segmentation method based on the distribution function of an image", Computer vision, graphics, and image processing, Vol. 29, pp. 47-59.

BOW, S., 1984, "Pattern recognition - Applications to large data-set problems", Marcel Dekker.

BROWN, C.M., 1983, "Inherent bias and noise in the Hough transform", IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. PAMI-5, No. 5, Sep., pp. 493-505.

BROWN, G.S., 1965, "The yield of clearwood from pruning: some results with radiata pine", Commonwealth Forestry Review, Vol. 44, No. 3, pp. 197-221.

BRUNELLE, N.R., HIGGINS, F.P., 1986, "Line scan vision system", AT&T Technical Journal, Jul/Aug, Vol. 65, No. 4, pp. 58-65.

CHIN, R.T., HARLOW, C.A., 1982, "Automated visual inspection: a survey", IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. PAMI-4, No. 6, Nov., pp. 557-573.

CONNERS, R.W., HARLOW, C.A., 1978, "Equal probability quantising and texture analysis of radiographic images", Computer Graphics and Image Processing, Vol. 8, pp. 447-463.

CONNERS, R.W., HARLOW, C.A., 1980, "A theoretical comparison of texture algorithms", IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. PAMI-2, No. 3, May., pp. 204-222.

CONNERS, R.W., McMILLIN, C.W., LIN, K., VASQUEZ-ESPINOSA, R.E., 1983, "Identifying and locating surface defects in wood: Part of an automated lumber processing system", IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. PAMI-5, No. 6, Nov., pp. 573-583.

CONNERS, R.W., McMILLIN, C.W., VASQUEZ-ESPIOSA, R.E., 1984, "A prototype system for locating and identifying surface defects in wood", Proc. 7th Int. Conf. Pattern Recognition (ICPR), Montreal, P.Q., Canada, Aug., pp. 416-419.

CULLINAN, M.C., 1984, "Automated visual inspection for industrial quality assessment and control: local experiences", Proceedings National Conference and Exhibition on Robotics, Melbourne, Australia, Aug. , pp. 114-117.

CHIN, R.T., 1988, "Automated visual inspection: 1981 to 1987", Computer vision, graphics, and image processing, Vol. 41, pp. 346-381.

DAVENEL, A., GUIZARD, CH., LABARRE, T., SEVILA, F., 1988, "Automatic detection of surface defects on fruit by using a vision system", Journal of Agricultural Engineering Research, Vol. 41, pp. 1-9.

DINSTEIN, I., FONG, A.C., NI, L.M., WONG, K.Y., 1984, "Fast discrimination between homogeneous and textured regions", Pattern Recognition: Proceedings of the 7th International Conference, Vol. 1, pp. 361-363.

DYER, C.R., 1983, "Gauge inspection using Hough transforms", IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. PAMI-5, No. 6, Nov., pp. 621-623.

DYKES, G.W., 1985, "Automated inspection of food jars for glass fragments", Vision '85 Conference Proceedings, pp. 621-630.

FOLEY, D., 1971, "The probability of error on the design set as a function of the sample size and feature size",Ph. D. dissertion, Syracuse University, June. Available as Rome Air Development Centre Report, RADC-TR-71-171. (in Meisel, 1972, p. 35)

FU, K.S., 1983, "Pictorial pattern recognition for industrial inspection", in HARALICK, R. (Ed.): "Pictorial data analysis", Springer-Verlag, pp. 335-349.

GEPP, B.C., 1986, in WALLACE, H.R., "The ecology of the forests and woodlands of South Australia", Government Printer, South Australia.

GONZALEZ, R.C., WINTZ, P., 1987, "Digital image processing", Addison-Wesley.

HARA, Y., AKIYAMA, N., KARASAKI, K., 1983, "Automatic inspection system for printed circuit boards", IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. PAMI-5, No. 6, Nov., pp. 623-630.

HARALICK, R.M., SHANMUGAM, K., 1973(a), "Computer classification of reservoir sandstones", IEEE Trans. Geoscience Electronics, Vol. GE-11, Oct., pp 171-177.

HARALICK, R.M., SHANMUGAM, K., DINSTEIN, I., 1973(b), "Textural features for image classification", IEEE Trans. on Systems, Man, and Cybernetics, Vol. SMC-3, No. 6, Nov., pp 610-21.

HARALICK, R.M., 1983, "Image texture survey", Fundamentals in Computer Vision, Ed. O.D.Faugeras, pp. 145-172.

HO, Y., AGRAWALA, A.K., 1968, "On pattern classification algorithms - introduction and survey", Proceedings of the IEEE, Vol. 56, Dec., pp. 2101-2114.

JOSEPH, A., SHARMA, A.K., 1985, "Visual grading of timber with a computer", CIVIL-COMP 85, Proceedings of the second international conference on civil and structural computing, Vol. 1, pp. 239-246.

KAPUR, J.N., SAHOO, P.K., WONG, A.K.C., 1985, " A new method for grey-level picture thresholding using the entropy of the histogram", Computer vision, graphics, and image processing, Vol. 29, pp. 273-285.

KING, E.A., 1978, "Laser scanning", Proceedings 4th Nondestructive Testing of Wood Symposium, Vancouver, WA, Aug. 28-30, pp. 15-22.

KLAUS, B., HORN, P., 1986, Robot Vision, MIT Press, pp. 46-61.

LAVERY, P.B., 1986, "Pinus radiata - an introduction to the species", Plantaion Forestry with Pinus Radiata, Review Papers, Paper No. 12, School of Forestry, University of Canterbury, Christchurch, New Zealand, pp. 5-21.

LIN, K., 1983, "Application of image analysis to lumber inspection", ACM Computer Science Conference and SIGCSE Symposium, New York, NY, USA, Feb., p. 74.

MATHEWS, P.C., BEECH, B.H., 1976, "Method and apparatus for detecting timber defects", US Patent 3,976,384, Aug. 24.

MATHEWS, P.C., BEECH, B.H., 1977, "Improvements in or relating to optical detection apparatus", British Patent 1,488,841, Oct. 12.

MEISEL, W.S., 1972, "Computer-oriented approaches to pattern recognition", Academic Press.

MUELLER, P.A., HEBERT, R.T., 1976, "Increasing production yields with a laser scanner lumber inspection system", IEEE Optical Society of America Conference, San Diego, California, USA, 25-27 May, p. 6.

NAKAGAWA, Y., NINOMIYA, T., 1984, "The structured light method for inspection of solder joints and assembly robot vision systems", in Robotics Research, Ed. M. Brady and R. Paul, MIT Press, pp. 355-369.

NILSSON, N.J., 1965, "Learning machines", McGraw-Hill.

OCTEK Inc, 1986, "An automated fish inspection system", Robotics Engineering, Jun., p. 19.

OKAWA, Y., 1984, "Automatic inspection of the surface defects of cast metals", Computer vision, graphics, and image processing, Vol. 25, pp. 89-112.

PACKER, M.E., 1979, "The application of image processing to the detection of corrosion by radiography", Defence Science and Technology Organisation, Aeronautical Research Labs., Materials Report 109, Feb.

PAI, A.L., LEE, K., PALMER, K.L., SELVIDGE, D.G., 1986, "Automated visual inspection of aircraft engine combustor assemblies", Proc. IEEE Int. Conf. on Robotics and Automation, pp. 1919-1924.

PAI, A.L., PATEL,.K.U., LAI, K.K., 1988, "A feasibility study of IC chip wire bond inspection", Journal of Robotic Systems, Vol. 5, No. 2, pp. 147-179.

PATRICK, E.A., 1972, "Fundamentals of pattern recognition", Prentice-Hall.

PAUL, D., HATTICH, W., NILL, W., TATARI, S., WINKLER, G., 1988, "VISTA: Visual Interpretation System for Technical Applications - architecture and use", IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. PAMI-10, No. 3, May, pp. 399-407.

PETERSON, V., EVEN, J., 1985, "A microcomputer based binary vision system", International Journal of Applied Engineering Education, Vol. 1, No. 6, pp. 405-413.

POELZLEITNER, W., 1986, "A Hough transform method to segment images of wooden boards", Proc. 8th Int. Conf. Pattern Recognition (ICPR), Paris, France, Oct., pp. 262-264.

RAY, R., 1988, "Automated visual inspection of solder bumps", AT&T Technical Journal, Mar/Apr, Vol. 67, No. 2, pp. 47-60.

ROUGHANA, E.J., 1988, "The future of sawn timber in the green triangle region", Appita Mini Conference, International Motel, Mount Gambier, Nov.

SAHOO, P.K., SOLTANI, S., WONG, A.K.C., CHEN, Y.C., 1988, "A survey of thresholding techniques", Computer vision, graphics, and image processing, Vol. 41, pp. 233-260.

SCHUERMANN, J., 1977, "Polynomklassifikatoren fuer die Zeichenerkennung", Munich, FRG; Vienna, Austria.

SCOTT, C.W., 1960, "Pinus Radiata", Food and Agriculture Organisation of the UN, Rome.

SEBESTYEN, G., 1962, "Decision-making processes in pattern recognition", Macmillan, New York.

SHAPIRO, S.D., 1978, "Generalisation of the Hough transform for curve detection in noisy digital pictures", Proceedings 4th Int. Conf. Pattern Recognition, Kyoto, Japan, pp. 70-80.

SIEW, L.H., HODGSON, R.M., WOOD, E.J., 1988, "Texture measures for carpet wear assessment", IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. PAMI-10, No. 1, Jan., pp. 92-105.

SZYMANI, R., McDONALD, K.A., 1981, "Defect detection in lumber: state of the art", Forest Poducts Journal, Vol. 31, Nov, pp. 34-43.

TATARI, S., HATTICH, W., 1987(a), "Automatic recognition of defects in wood", SPIE Proc. Adv. Image Processing, Vol. 804, pp. 229-236.

TATARI, S., 1987(b), "Selection and parametrisation of texture analysis methods for the automation of industrial visual inspection tasks", Applied Digital Image Processing X, SPIE Proc. 31st Annual Int. Tech. Symp., San Diego, CA, Aug., pp. 86-94.

THIBADEAU, R., 1986, "Report on the CMU-Westinghouse printed wiring board inspection device: PWBIS II", SPIE Proc. Automated Inspection and Measurement, Vol. 730, Oct., pp. 174-186.

TSAI, W., 1985, "Moment-preserving thresholding: A new approach", Computer vision, graphics, and image processing, Vol. 29, pp. 377-393.

VAN GOOL, L., DEWAELE, P., OOSTERLINCK, A., 1985, "Texture analysis anno 1983", Computer vision, graphics, and image processing, Vol. 29, pp. 336-357.

WALLACE, A.M., 1988, "Industrial applications of computer vision since 1982", IEE proceedings, Vol. 135, Pt. E, No. 3, May, pp. 117-136.

WARD, M.R., ROSSOL, L., HOLLAND, S.W., 1979, "CONSIGHT: A practical vision based robot guidance system", GM Research Publication, GMR-2912.

WESZKA, J.S., DYER, C.R., ROSENFELD, A., 1976, "A comparative study of texture measures for terrain classification", IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-6, No. 4, Apr., pp. 269-285.

WESZKA, J.S., ROSENFELD, A., 1978, "Threshold evaluation techniques", IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-8, No. 8, Aug., pp. 622-629.

ZAHNISER, D.J., 1983, "Automation of pap smear analysis: a review and status report", Pictorial Data Analysis, Ed. R.M. Haralick, NATO ASI Series, Vol. F4, Springer-Verlag, pp. 265-294.